

CIS Amazon Linux 2023 Benchmark

v1.0.0 - 06-26-2023

Terms of Use

Please see the below link for our current terms of use:

<https://www.cisecurity.org/cis-securesuite/cis-securesuite-membership-terms-of-use/>

Table of Contents

Terms of Use	1
Table of Contents	2
Overview	9
Intended Audience.....	9
Consensus Guidance	10
Typographical Conventions.....	11
Recommendation Definitions.....	12
Title.....	12
Assessment Status.....	12
Automated	12
Manual.....	12
Profile	12
Description.....	12
Rationale Statement	12
Impact Statement.....	13
Audit Procedure.....	13
Remediation Procedure.....	13
Default Value.....	13
References	13
CIS Critical Security Controls® (CIS Controls®)	13
Additional Information.....	13
Profile Definitions	14
Acknowledgements	15
Recommendations	16
1 Initial Setup	16
1.1 Filesystem Configuration.....	17
1.1.1 Configure Filesystem Kernel Modules.....	18
1.1.1.1 Ensure mounting of squashfs filesystems is disabled (Automated)	19
1.1.1.2 Ensure mounting of udf filesystems is disabled (Automated).....	24
1.1.1.3 Ensure mounting of cramfs filesystems is disabled (Automated)	29
1.1.1.4 Ensure mounting of freevxfs filesystems is disabled (Automated)	34
1.1.1.5 Ensure mounting of jffs2 filesystems is disabled (Automated)	39
1.1.1.6 Ensure mounting of hfs filesystems is disabled (Automated)	44
1.1.1.7 Ensure mounting of hfsplus filesystems is disabled (Automated)	49
1.1.2 Configure /tmp.....	54
1.1.2.1 Ensure /tmp is a separate partition (Automated).....	55
1.1.2.2 Ensure nodev option set on /tmp partition (Automated)	58

1.1.2.3 Ensure noexec option set on /tmp partition (Automated)	60
1.1.2.4 Ensure nosuid option set on /tmp partition (Automated)	62
1.1.3 Configure /var	64
1.1.3.1 Ensure separate partition exists for /var (Automated)	65
1.1.3.2 Ensure nodev option set on /var partition (Automated)	68
1.1.3.3 Ensure nosuid option set on /var partition (Automated)	70
1.1.4 Configure /var/tmp	72
1.1.4.1 Ensure separate partition exists for /var/tmp (Automated)	73
1.1.4.2 Ensure noexec option set on /var/tmp partition (Automated)	76
1.1.4.3 Ensure nosuid option set on /var/tmp partition (Automated)	78
1.1.4.4 Ensure nodev option set on /var/tmp partition (Automated)	80
1.1.5 Configure /var/log	82
1.1.5.1 Ensure separate partition exists for /var/log (Automated)	83
1.1.5.2 Ensure nodev option set on /var/log partition (Automated)	85
1.1.5.3 Ensure noexec option set on /var/log partition (Automated)	87
1.1.5.4 Ensure nosuid option set on /var/log partition (Automated)	89
1.1.6 Configure /var/log/audit	91
1.1.6.1 Ensure separate partition exists for /var/log/audit (Automated)	92
1.1.6.2 Ensure noexec option set on /var/log/audit partition (Automated)	94
1.1.6.3 Ensure nodev option set on /var/log/audit partition (Automated)	96
1.1.6.4 Ensure nosuid option set on /var/log/audit partition (Automated)	98
1.1.7 Configure /home	100
1.1.7.1 Ensure separate partition exists for /home (Automated)	101
1.1.7.2 Ensure nodev option set on /home partition (Automated)	104
1.1.7.3 Ensure nosuid option set on /home partition (Automated)	106
1.1.8 Configure /dev/shm	108
1.1.8.1 Ensure /dev/shm is a separate partition (Automated)	109
1.1.8.2 Ensure nodev option set on /dev/shm partition (Automated)	111
1.1.8.3 Ensure noexec option set on /dev/shm partition (Automated)	113
1.1.8.4 Ensure nosuid option set on /dev/shm partition (Automated)	115
1.1.9 Ensure usb-storage is disabled (Automated)	117
1.2 Configure Software and Patch Management	121
1.2.1 Ensure GPG keys are configured (Manual)	122
1.2.2 Ensure gpgcheck is globally activated (Automated)	125
1.2.3 Ensure package manager repositories are configured (Manual)	127
1.2.4 Ensure repo_gpgcheck is globally activated (Manual)	129
1.3 Filesystem Integrity Checking	131
1.3.1 Ensure AIDE is installed (Automated)	132
1.3.2 Ensure filesystem integrity is regularly checked (Automated)	134
1.3.3 Ensure cryptographic mechanisms are used to protect the integrity of audit tools (Automated)	137
1.4 Secure Boot Settings	139
1.4.1 Ensure permissions on bootloader config are configured (Automated)	140
1.5 Additional Process Hardening	144
1.5.1 Ensure address space layout randomization (ASLR) is enabled (Automated)	145
1.5.2 Ensure ptrace_scope is restricted (Automated)	149
1.5.3 Ensure core dump storage is disabled (Automated)	153
1.5.4 Ensure core dump backtraces are disabled (Automated)	155
1.6 Mandatory Access Control	157
1.6.1 Configure SELinux	158
1.6.1.1 Ensure SELinux is installed (Automated)	160
1.6.1.2 Ensure SELinux is not disabled in bootloader configuration (Automated)	162

1.6.1.3 Ensure SELinux policy is configured (Automated)	165
1.6.1.4 Ensure the SELinux mode is not disabled (Automated)	167
1.6.1.5 Ensure the SELinux mode is enforcing (Automated)	170
1.6.1.6 Ensure no unconfined services exist (Manual)	173
1.6.1.7 Ensure SETroubleshoot is not installed (Automated)	175
1.6.1.8 Ensure the MCS Translation Service (mcstrans) is not installed (Automated)	177
1.7 Command Line Warning Banners	179
1.7.1 Ensure message of the day is configured properly (Automated)	180
1.7.2 Ensure local login warning banner is configured properly (Automated)	182
1.7.3 Ensure remote login warning banner is configured properly (Automated)	184
1.7.4 Ensure permissions on /etc/motd are configured (Automated)	186
1.7.5 Ensure permissions on /etc/issue are configured (Automated)	188
1.7.6 Ensure permissions on /etc/issue.net are configured (Automated)	190
1.8 Ensure updates, patches, and additional security software are installed (Manual)	192
1.9 Ensure system-wide crypto policy is not legacy (Automated)	195
2 Services	197
2.1 Configure Time Synchronization	198
2.1.1 Ensure time synchronization is in use (Automated)	199
2.1.2 Ensure chrony is configured (Automated)	201
2.2 Configure Special Purpose Services	203
2.2.1 Ensure xorg-x11-server-common is not installed (Automated)	204
2.2.2 Ensure avahi is not installed (Automated)	206
2.2.3 Ensure a print server is not installed (Automated)	208
2.2.4 Ensure a dhcp server is not installed (Automated)	210
2.2.5 Ensure a dns server is not installed (Automated)	212
2.2.6 Ensure an ftp server is not installed (Automated)	214
2.2.7 Ensure a tftp server is not installed (Automated)	216
2.2.8 Ensure a web server is not installed (Automated)	218
2.2.9 Ensure IMAP and POP3 server is not installed (Automated)	220
2.2.10 Ensure Samba is not installed (Automated)	222
2.2.11 Ensure HTTP Proxy Server is not installed (Automated)	224
2.2.12 Ensure net-snmp is not installed or the snmpd service is not enabled (Automated)	226
2.2.13 Ensure telnet-server is not installed (Automated)	229
2.2.14 Ensure dnsmasq is not installed (Automated)	231
2.2.15 Ensure mail transfer agent is configured for local-only mode (Automated)	232
2.2.16 Ensure nfs-utils is not installed or the nfs-server service is masked (Automated)	234
2.2.17 Ensure rpcbind is not installed or the rpcbind services are masked (Automated)	237
2.2.18 Ensure rsync-daemon is not installed or the rsyncd service is masked (Automated)	240
2.3 Service Clients	242
2.3.1 Ensure telnet client is not installed (Automated)	243
2.3.2 Ensure LDAP client is not installed (Automated)	245
2.3.3 Ensure FTP client is not installed (Automated)	247
2.4 Ensure nonessential services listening on the system are removed or masked (Manual)	249
3 Network Configuration	251
3.1 Disable unused network protocols and devices	252
3.1.1 Ensure IPv6 status is identified (Manual)	253
3.1.2 Ensure DCCP is disabled (Automated)	255
3.1.3 Ensure SCTP is disabled (Automated)	260
3.1.4 Ensure RDS is disabled (Automated)	265
3.1.5 Ensure TIPC is disabled (Automated)	270
3.2 Configure Network Parameters (Host Only)	275

3.2.1 Ensure IP forwarding is disabled (Automated)	276
3.2.2 Ensure packet redirect sending is disabled (Automated)	280
3.3 Configure Network Parameters (Host and Router)	284
3.3.1 Ensure source routed packets are not accepted (Automated)	285
3.3.2 Ensure ICMP redirects are not accepted (Automated)	290
3.3.3 Ensure secure ICMP redirects are not accepted (Automated)	294
3.3.4 Ensure suspicious packets are logged (Automated)	298
3.3.5 Ensure broadcast ICMP requests are ignored (Automated)	302
3.3.6 Ensure bogus ICMP responses are ignored (Automated)	306
3.3.7 Ensure Reverse Path Filtering is enabled (Automated)	310
3.3.8 Ensure TCP SYN Cookies is enabled (Automated)	314
3.3.9 Ensure IPv6 router advertisements are not accepted (Automated)	318
3.4 Configure Host Based Firewall	322
3.4.1 Configure a firewall utility	323
3.4.1.1 Ensure nftables is installed (Automated)	324
3.4.1.2 Ensure a single firewall configuration utility is in use (Automated)	326
3.4.2 Configure firewall rules	330
3.4.2.1 Ensure firewalld default zone is set (Automated)	334
3.4.2.2 Ensure at least one nftables table exists (Automated)	337
3.4.2.3 Ensure nftables base chains exist (Automated)	339
3.4.2.4 Ensure host based firewall loopback traffic is configured (Automated)	341
3.4.2.5 Ensure firewalld drops unnecessary services and ports (Manual)	346
3.4.2.6 Ensure nftables established connections are configured (Manual)	348
3.4.2.7 Ensure nftables default deny firewall policy (Automated)	350
4 Access, Authentication and Authorization	352
4.1 Configure time-based job schedulers	353
4.1.1 Ensure cron daemon is installed and enabled (Automated)	354
4.1.2 Ensure permissions on /etc/crontab are configured (Automated)	356
4.1.3 Ensure permissions on /etc/cron.hourly are configured (Automated)	358
4.1.4 Ensure permissions on /etc/cron.daily are configured (Automated)	360
4.1.5 Ensure permissions on /etc/cron.weekly are configured (Automated)	362
4.1.6 Ensure permissions on /etc/cron.monthly are configured (Automated)	364
4.1.7 Ensure permissions on /etc/cron.d are configured (Automated)	366
4.1.8 Ensure cron is restricted to authorized users (Automated)	368
4.1.9 Ensure at is restricted to authorized users (Automated)	370
4.2 Configure SSH Server	373
4.2.1 Ensure permissions on /etc/ssh/ssh_config are configured (Automated)	374
4.2.2 Ensure permissions on SSH private host key files are configured (Automated)	376
4.2.3 Ensure permissions on SSH public host key files are configured (Automated)	380
4.2.4 Ensure SSH access is limited (Automated)	384
4.2.5 Ensure SSH LogLevel is appropriate (Automated)	387
4.2.6 Ensure SSH PAM is enabled (Automated)	390
4.2.7 Ensure SSH root login is disabled (Automated)	392
4.2.8 Ensure SSH HostbasedAuthentication is disabled (Automated)	394
4.2.9 Ensure SSH PermitEmptyPasswords is disabled (Automated)	396
4.2.10 Ensure SSH PermitUserEnvironment is disabled (Automated)	398
4.2.11 Ensure SSH IgnoreRhosts is enabled (Automated)	400
4.2.12 Ensure SSH X11 forwarding is disabled (Automated)	402
4.2.13 Ensure SSH AllowTcpForwarding is disabled (Automated)	404
4.2.14 Ensure system-wide crypto policy is not over-ridden (Automated)	407
4.2.15 Ensure SSH warning banner is configured (Automated)	408

4.2.16 Ensure SSH MaxAuthTries is set to 4 or less (Automated)	409
4.2.17 Ensure SSH MaxStartups is configured (Automated)	411
4.2.18 Ensure SSH MaxSessions is set to 10 or less (Automated)	414
4.2.19 Ensure SSH LoginGraceTime is set to one minute or less (Automated)	416
4.2.20 Ensure SSH Idle Timeout Interval is configured (Automated)	418
4.3 Configure privilege escalation	421
4.3.1 Ensure sudo is installed (Automated)	422
4.3.2 Ensure sudo commands use pty (Automated)	424
4.3.3 Ensure sudo log file exists (Automated)	426
4.3.4 Ensure re-authentication for privilege escalation is not disabled globally (Automated)	428
4.3.5 Ensure sudo authentication timeout is configured correctly (Automated)	430
4.3.6 Ensure access to the su command is restricted (Automated)	432
4.4 Configure authselect	434
4.4.1 Ensure custom authselect profile is used (Manual)	435
4.4.2 Ensure authselect includes with-faillock (Automated)	438
4.5 Configure PAM	440
4.5.1 Ensure password creation requirements are configured (Automated)	441
4.5.2 Ensure lockout for failed password attempts is configured (Automated)	445
4.5.3 Ensure password reuse is limited (Automated)	448
4.5.4 Ensure password hashing algorithm is SHA-512 (Automated)	451
4.6 User Accounts and Environment	455
4.6.1 Set Shadow Password Suite Parameters	456
4.6.1.1 Ensure password expiration is 365 days or less (Automated)	457
4.6.1.2 Ensure minimum days between password changes is configured (Automated)	459
4.6.1.3 Ensure password expiration warning days is 7 or more (Automated)	461
4.6.1.4 Ensure inactive password lock is 30 days or less (Automated)	463
4.6.1.5 Ensure all users last password change date is in the past (Automated)	465
4.6.2 Ensure system accounts are secured (Automated)	467
4.6.3 Ensure default user shell timeout is 900 seconds or less (Automated)	470
4.6.4 Ensure default group for the root account is GID 0 (Automated)	474
4.6.5 Ensure default user umask is 027 or more restrictive (Automated)	476
4.6.6 Ensure root password is set (Manual)	481
5 Logging and Auditing	483
5.1 Configure Logging	484
5.1.1 Configure rsyslog	485
5.1.1.1 Ensure rsyslog is installed (Manual)	486
5.1.1.2 Ensure rsyslog service is enabled (Manual)	488
5.1.1.3 Ensure journald is configured to send logs to rsyslog (Manual)	490
5.1.1.4 Ensure rsyslog default file permissions are configured (Automated)	493
5.1.1.5 Ensure logging is configured (Manual)	496
5.1.1.6 Ensure rsyslog is configured to send logs to a remote log host (Manual)	498
5.1.1.7 Ensure rsyslog is not configured to receive logs from a remote client (Automated)	500
5.1.2 Configure journald	503
5.1.2.1.1 Ensure systemd-journal-remote is installed (Manual)	505
5.1.2.1.2 Ensure systemd-journal-remote is configured (Manual)	507
5.1.2.1.3 Ensure systemd-journal-remote is enabled (Manual)	509
5.1.2.1.4 Ensure journald is not configured to receive logs from a remote client (Automated)	511
5.1.2.2 Ensure journald service is enabled (Automated)	513
5.1.2.3 Ensure journald is configured to compress large log files (Automated)	515
5.1.2.4 Ensure journald is configured to write logfiles to persistent disk (Automated)	517
5.1.2.5 Ensure journald is not configured to send logs to rsyslog (Manual)	519

5.1.2.6 Ensure journald log rotation is configured per site policy (Manual)	521
5.1.2.7 Ensure journald default file permissions configured (Manual)	523
5.1.3 Ensure all logfiles have appropriate permissions and ownership (Automated)	525
5.2 Configure System Accounting (auditd)	531
5.2.1 Ensure auditing is enabled	533
5.2.1.1 Ensure auditd is installed (Automated).....	534
5.2.1.2 Ensure auditing for processes that start prior to auditd is enabled (Automated)	536
5.2.1.3 Ensure audit_backlog_limit is sufficient (Automated)	538
5.2.1.4 Ensure auditd service is enabled (Automated).....	540
5.2.2 Configure Data Retention	542
5.2.2.1 Ensure audit log storage size is configured (Automated)	543
5.2.2.2 Ensure audit logs are not automatically deleted (Automated)	545
5.2.2.3 Ensure system is disabled when audit logs are full (Automated)	547
5.2.3 Configure auditd rules	550
5.2.3.1 Ensure changes to system administration scope (sudoers) is collected (Automated).....	551
5.2.3.2 Ensure actions as another user are always logged (Automated)	554
5.2.3.3 Ensure events that modify the sudo log file are collected (Automated).....	558
5.2.3.4 Ensure events that modify date and time information are collected (Automated)	562
5.2.3.5 Ensure events that modify the system's network environment are collected (Automated).....	566
5.2.3.6 Ensure use of privileged commands are collected (Automated)	570
5.2.3.7 Ensure unsuccessful file access attempts are collected (Automated).....	574
5.2.3.8 Ensure events that modify user/group information are collected (Automated)	578
5.2.3.9 Ensure discretionary access control permission modification events are collected (Automated) ..	582
5.2.3.10 Ensure successful file system mounts are collected (Automated)	587
5.2.3.11 Ensure session initiation information is collected (Automated)	591
5.2.3.12 Ensure login and logout events are collected (Automated)	595
5.2.3.13 Ensure file deletion events by users are collected (Automated).....	598
5.2.3.14 Ensure events that modify the system's Mandatory Access Controls are collected (Automated)602	
5.2.3.15 Ensure successful and unsuccessful attempts to use the chcon command are recorded (Automated)	606
5.2.3.16 Ensure successful and unsuccessful attempts to use the setfacl command are recorded (Automated)	610
5.2.3.17 Ensure successful and unsuccessful attempts to use the chacl command are recorded (Automated)	614
5.2.3.18 Ensure successful and unsuccessful attempts to use the usermod command are recorded (Automated)	618
5.2.3.19 Ensure kernel module loading unloading and modification is collected (Automated).....	622
5.2.3.20 Ensure the audit configuration is immutable (Automated).....	627
5.2.3.21 Ensure the running and on disk configuration is the same (Manual).....	629
5.2.4 Configure auditd file access	631
5.2.4.1 Ensure audit log files are mode 0640 or less permissive (Automated)	632
5.2.4.2 Ensure only authorized users own audit log files (Automated).....	634
5.2.4.3 Ensure only authorized groups are assigned ownership of audit log files (Automated)	636
5.2.4.4 Ensure the audit log directory is 0750 or more restrictive (Automated).....	638
5.2.4.5 Ensure audit configuration files are 640 or more restrictive (Automated).....	640
5.2.4.6 Ensure audit configuration files are owned by root (Automated).....	642
5.2.4.7 Ensure audit configuration files belong to group root (Automated)	644
5.2.4.8 Ensure audit tools are 755 or more restrictive (Automated)	646
5.2.4.9 Ensure audit tools are owned by root (Automated)	648
5.2.4.10 Ensure audit tools belong to group root (Automated).....	650
5.3 Ensure logrotate is configured (Manual)	652
6 System Maintenance	654

6.1 System File Permissions	655
6.1.1 Ensure permissions on /etc/passwd are configured (Automated)	656
6.1.2 Ensure permissions on /etc/passwd are configured (Automated)	658
6.1.3 Ensure permissions on /etc/passwd- are configured (Automated)	660
6.1.4 Ensure permissions on /etc/group are configured (Automated)	662
6.1.5 Ensure permissions on /etc/group- are configured (Automated)	664
6.1.6 Ensure permissions on /etc/shadow are configured (Automated)	666
6.1.7 Ensure permissions on /etc/shadow- are configured (Automated)	668
6.1.8 Ensure permissions on /etc/gshadow are configured (Automated)	670
6.1.9 Ensure permissions on /etc/gshadow- are configured (Automated)	672
6.1.10 Audit system file permissions (Manual)	674
6.1.11 Ensure world writable files and directories are secured (Automated)	677
6.1.12 Ensure no unowned or ungrouped files or directories exist (Automated)	682
6.1.13 Ensure SUID and SGID files are reviewed (Manual)	685
6.2 Local User and Group Settings	688
6.2.1 Ensure accounts in /etc/passwd use shadowed passwords (Automated)	689
6.2.2 Ensure /etc/shadow password fields are not empty (Automated)	691
6.2.3 Ensure all groups in /etc/passwd exist in /etc/group (Automated)	693
6.2.4 Ensure no duplicate UIDs exist (Automated)	694
6.2.5 Ensure no duplicate GIDs exist (Automated)	695
6.2.6 Ensure no duplicate user names exist (Automated)	696
6.2.7 Ensure no duplicate group names exist (Automated)	697
6.2.8 Ensure root PATH Integrity (Automated)	698
6.2.9 Ensure root is the only UID 0 account (Automated)	700
6.2.10 Ensure local interactive user home directories are configured (Automated)	701
6.2.11 Ensure local interactive user dot files access is configured (Automated)	706
Appendix: Summary Table	712
Appendix: Change History	779

Overview

All CIS Benchmarks focus on technical configuration settings used to maintain and/or increase the security of the addressed technology, and they should be used in **conjunction** with other essential cyber hygiene tasks like:

- Monitoring the base operating system for vulnerabilities and quickly updating with the latest security patches
- Monitoring applications and libraries for vulnerabilities and quickly updating with the latest security patches

In the end, the CIS Benchmarks are designed as a key **component** of a comprehensive cybersecurity program.

This document provides prescriptive guidance for establishing a secure configuration posture for Amazon Linux 2023 running on x86_64 platforms.

The guidance within broadly assumes that operations are being performed as the `root` user, and executed under the default Bash version for the applicable distribution.

Operations performed using `sudo` instead of the `root` user, or executed under another shell, may produce unexpected results, or fail to make the intended changes to the system. Non-root users may not be able to access certain areas of the system, especially after remediation has been performed. It is advisable to verify `root` users path integrity and the integrity of any programs being run prior to execution of commands and scripts included in this benchmark.

The default prompt for the `root` user is `#`, and as such all sample commands will have `#` as an additional indication that it is to be executed as `root`.

To obtain the latest version of this guide, please visit <http://workbench.cisecurity.org>. If you have questions, comments, or have identified ways to improve this guide, please write us at feedback@cisecurity.org.

Intended Audience

This benchmark is intended for system and application administrators, security specialists, auditors, help desk, and platform deployment personnel who plan to develop, deploy, assess, or secure solutions that incorporate Amazon Linux 2023 on x86_64 platforms.

Consensus Guidance

This CIS Benchmark was created using a consensus review process comprised of a global community of subject matter experts. The process combines real world experience with data-based information to create technology specific guidance to assist users to secure their environments. Consensus participants provide perspective from a diverse set of backgrounds including consulting, software development, audit and compliance, security research, operations, government, and legal.

Each CIS Benchmark undergoes two phases of consensus review. The first phase occurs during initial Benchmark development. During this phase, subject matter experts convene to discuss, create, and test working drafts of the Benchmark. This discussion occurs until consensus has been reached on Benchmark recommendations. The second phase begins after the Benchmark has been published. During this phase, all feedback provided by the Internet community is reviewed by the consensus team for incorporation in the Benchmark. If you are interested in participating in the consensus process, please visit <https://workbench.cisecurity.org/>.

Typographical Conventions

The following typographical conventions are used throughout this guide:

Convention	Meaning
<code>Stylized Monospace font</code>	Used for blocks of code, command, and script examples. Text should be interpreted exactly as presented.
<code>Monospace font</code>	Used for inline code, commands, or examples. Text should be interpreted exactly as presented.
<i><italic font in brackets></i>	Italic texts set in angle brackets denote a variable requiring substitution for a real value.
<i>Italic font</i>	Used to denote the title of a book, article, or other publication.
Note	Additional information or caveats

Recommendation Definitions

The following defines the various components included in a CIS recommendation as applicable. If any of the components are not applicable it will be noted or the component will not be included in the recommendation.

Title

Concise description for the recommendation's intended configuration.

Assessment Status

An assessment status is included for every recommendation. The assessment status indicates whether the given recommendation can be automated or requires manual steps to implement. Both statuses are equally important and are determined and supported as defined below:

Automated

Represents recommendations for which assessment of a technical control can be fully automated and validated to a pass/fail state. Recommendations will include the necessary information to implement automation.

Manual

Represents recommendations for which assessment of a technical control cannot be fully automated and requires all or some manual steps to validate that the configured state is set as expected. The expected state can vary depending on the environment.

Profile

A collection of recommendations for securing a technology or a supporting platform. Most benchmarks include at least a Level 1 and Level 2 Profile. Level 2 extends Level 1 recommendations and is not a standalone profile. The Profile Definitions section in the benchmark provides the definitions as they pertain to the recommendations included for the technology.

Description

Detailed information pertaining to the setting with which the recommendation is concerned. In some cases, the description will include the recommended value.

Rationale Statement

Detailed reasoning for the recommendation to provide the user a clear and concise understanding on the importance of the recommendation.

Impact Statement

Any security, functionality, or operational consequences that can result from following the recommendation.

Audit Procedure

Systematic instructions for determining if the target system complies with the recommendation

Remediation Procedure

Systematic instructions for applying recommendations to the target system to bring it into compliance according to the recommendation.

Default Value

Default value for the given setting in this recommendation, if known. If not known, either not configured or not defined will be applied.

References

Additional documentation relative to the recommendation.

CIS Critical Security Controls® (CIS Controls®)

The mapping between a recommendation and the CIS Controls is organized by CIS Controls version, Safeguard, and Implementation Group (IG). The Benchmark in its entirety addresses the CIS Controls safeguards of (v7) "5.1 - Establish Secure Configurations" and (v8) "4.1 - Establish and Maintain a Secure Configuration Process" so individual recommendations will not be mapped to these safeguards.

Additional Information

Supplementary information that does not correspond to any other field but may be useful to the user.

Profile Definitions

The following configuration profiles are defined by this Benchmark:

- **Level 1 - Server**

Items in this profile intend to:

- be practical and prudent;
- provide a clear security benefit; and
- not inhibit the utility of the technology beyond acceptable means.

This profile is intended for servers.

- **Level 2 - Server**

This profile extends the "Level 1 - Server" profile. Items in this profile exhibit one or more of the following characteristics:

- are intended for environments or use cases where security is paramount.
- acts as defense in depth measure.
- may negatively inhibit the utility or performance of the technology.

This profile is intended for servers.

Acknowledgements

This Benchmark exemplifies the great things a community of users, vendors, and subject matter experts can accomplish through consensus collaboration. The CIS community thanks the entire consensus team with special recognition to the following individuals who contributed greatly to the creation of this guide:

This benchmark is based upon previous Linux benchmarks published and would not be possible without the contributions provided over the history of all of these benchmarks. The CIS community thanks everyone who has contributed to the Linux benchmarks.

Contributor

Graham Eames
Simon John
Agustin Gonzalez
Tamas Tevesz
Beni Williamson

Editor

Jonathan Lewis Christopherson
Eric Pinnell
Justin Brown

Recommendations

1 Initial Setup

Items in this section are advised for all systems, but may be difficult or require extensive preparation after the initial setup of the system.

1.1 Filesystem Configuration

Directories that are used for system-wide functions can be further protected by placing them on separate partitions. This provides protection for resource exhaustion and enables the use of mounting options that are applicable to the directory's intended use. Users' data can be stored on separate partitions and have stricter mount options. A user partition is a filesystem that has been established for use by the users and does not contain software for system operations.

The recommendations in this section are easier to perform during initial system installation. If the system is already installed, it is recommended that a full backup be performed before repartitioning the system.

Note: If you are repartitioning a system that has already been installed (This may require the system to be in single-user mode):

- Mount the new partition to a temporary mountpoint e.g. `mount /dev/sda2 /mnt`
- Copy data from the original partition to the new partition. e.g. `cp /var/tmp/* /mnt`
- Verify that all data is present on the new partition. e.g. `ls -la /mnt`
- Unmount the new partition. e.g. `umount /mnt`
- Remove the data from the original directory that was in the old partition. e.g. `rm -Rf /var/tmp/*` Otherwise it will still consume space in the old partition that will be masked when the new filesystem is mounted.
- Mount the new partition to the desired mountpoint. e.g. `mount /dev/sda2 /var/tmp`
- Update `/etc/fstab` with the new mountpoint. e.g. `/dev/sda2 /var/tmp xfs defaults,rw,nosuid,nodev,noexec,relatime 0 0`

1.1.1 Configure Filesystem Kernel Modules

A number of uncommon filesystem types are supported under Linux. Removing support for unneeded filesystem types reduces the local attack surface of the system. If a filesystem type is not needed it should be disabled. Native Linux file systems are designed to ensure that built-in security controls function as expected. Non-native filesystems can lead to unexpected consequences to both the security and functionality of the system and should be used with caution. Many filesystems are created for niche use cases and are not maintained and supported as the operating systems are updated and patched. Users of non-native filesystems should ensure that there is attention and ongoing support for them, especially in light of frequent operating system changes.

Standard network connectivity and Internet access to cloud storage may make the use of non-standard filesystem formats to directly attach heterogeneous devices much less attractive.

Note: This should not be considered a comprehensive list of filesystems. You may wish to consider additions to those listed here for your environment. For the current available file system modules on the system see `/usr/lib/modules/$(uname -r)/kernel/fs`

Start up scripts

Kernel modules loaded directly via `insmod` will ignore what is configured in the relevant `/etc/modprobe.d/*.conf` files. If modules are still being loaded after a reboot whilst having the correctly configured `blacklist` and `install` command, check for `insmod` entries in start up scripts such as `.bashrc`.

You may also want to check `/lib/modprobe.d/`. Please note that this directory should not be used for user defined module loading. Ensure that all such entries resides in `/etc/modprobe.d/*.conf` files.

Return values

By using `/bin/false` as the command in disabling a particular module serves two purposes; to convey the meaning of the entry to the user and cause a non-zero return value. The latter can be tested for in scripts. Please note that `insmod` will ignore what is configured in the relevant `/etc/modprobe.d/*.conf` files. The preferred way to load modules is with `modprobe`.

1.1.1.1 Ensure mounting of squashfs filesystems is disabled (Automated)

Profile Applicability:

- Level 2 - Server

Description:

The `squashfs` filesystem type is a compressed read-only Linux filesystem embedded in small footprint systems. A `squashfs` image can be used without having to first decompress the image.

Rationale:

Removing support for unneeded filesystem types reduces the local attack surface of the system. If this filesystem type is not needed, disable it.

Impact:

As Snap packages utilizes `squashfs` as a compressed filesystem, disabling `squashfs` will cause Snap packages to fail.

`Snap` application packages of software are self-contained and work across a range of Linux distributions. This is unlike traditional Linux package management approaches, like APT or RPM, which require specifically adapted packages per Linux distribution on an application update and delay therefore application deployment from developers to their software's end-user. Snaps themselves have no dependency on any external store ("App store"), can be obtained from any source and can be therefore used for upstream software deployment.

Audit:

Run the following script to verify the `squashfs` module is disabled:

If the module is available in the running kernel:

- An entry including `/bin/true` or `/bin/false` exists in a file within the `/etc/modprobe.d/` directory
- The module is deny listed in a file within the `/etc/modprobe.d/` directory
- The module is not loaded in the kernel

If available in ANY installed kernel:

- The module is deny listed in a file within the `/etc/modprobe.d/` directory

If the kernel module is not available on the system, or pre-compiled into the kernel:

- No additional configuration is necessary

```

#!/usr/bin/env bash

{
    l_output="" l_output2="" l_output3="" l_dl="" # Unset output variables
    l_mname="squashfs" # set module name
    l_mtype="fs" # set module type
    l_searchloc="/lib/modprobe.d/*.conf /usr/local/lib/modprobe.d/*.conf /run/modprobe.d/*.conf
/etc/modprobe.d/*.conf"
    l_mpath="/lib/modules/**/kernel/$l_mtype"
    l_mpname="$(tr '-' '/' <<< "$l_mname")"
    l_mndir="$(tr '-' '/' <<< "$l_mname")"

    module_loadable_chk()
    {
        # Check if the module is currently loadable
        l_loadable="$(modprobe -n -v "$l_mname")"
        [ "$l_loadable" -gt "1" ] && l_loadable="$(grep -P --
"^(^\\h*install|\\b$l_mname)\\b" <<< "$l_loadable")"
        if grep -Pq -- '^\\h*install \\bin\\(true|false\\)' <<< "$l_loadable"; then
            l_output="$l_output\\n - module: \"$l_mname\" is not loadable: \"$l_loadable\""
        else
            l_output2="$l_output2\\n - module: \"$l_mname\" is loadable: \"$l_loadable\""
        fi
    }

    module_loaded_chk()
    {
        # Check if the module is currently loaded
        if ! lsmod | grep "$l_mname" > /dev/null 2>&1; then
            l_output="$l_output\\n - module: \"$l_mname\" is not loaded"
        else
            l_output2="$l_output2\\n - module: \"$l_mname\" is loaded"
        fi
    }

    module_deny_chk()
    {
        # Check if the module is deny listed
        l_dl="y"
        if modprobe --showconfig | grep -Pq -- '^\\h*blacklist\\h+'$l_mpname'"\\b'; then
            l_output="$l_output\\n - module: \"$l_mname\" is deny listed in: \"$(grep -Pls --
"^(\\h*blacklist\\h+$l_mname\\b" $l_searchloc)\""
        else
            l_output2="$l_output2\\n - module: \"$l_mname\" is not deny listed"
        fi
    }

    # Check if the module exists on the system
    for l_mdir in $l_mpath; do
        if [ -d "$l_mdir/$l_mndir" ] && [ -n "$(ls -A $l_mdir/$l_mndir)" ]; then
            l_output3="$l_output3\\n - \"$l_mdir\""
            [ "$l_dl" != "y" ] && module_deny_chk
            if [ "$l_mdir" = "/lib/modules/$(uname -r)/kernel/$l_mtype" ]; then
                module_loadable_chk
                module_loaded_chk
            fi
        else
            l_output="$l_output\\n - module: \"$l_mname\" doesn't exist in \"$l_mdir\""
        fi
    done

    # Report results. If no failures output in l_output2, we pass
    [ -n "$l_output3" ] && echo -e "\\n\\n -- INFO --\\n - module: \"$l_mname\" exists in:$l_output3"
    if [ -z "$l_output2" ]; then
        echo -e "\\n- Audit Result:\\n ** PASS **\\n$l_output\\n"
    else
        echo -e "\\n- Audit Result:\\n ** FAIL **\\n - Reason(s) for audit failure:\\n$l_output2\\n"
        [ -n "$l_output" ] && echo -e "\\n- Correctly set:\\n$l_output\\n"
    fi
}

```

Note: On operating systems where `squashfs` is pre-build into the kernel:

- This is considered an acceptable "passing" state
- The kernel **should not** be re-compiled to remove `squashfs`
- This audit will return as passing state with "module: "squashfs" doesn't exist in ..."

Remediation:

Run the following script to disable the `squashfs` module:

If the module is available in the running kernel:

- Create a file with `install squashfs /bin/false` in the `/etc/modprobe.d/` directory
- Create a file with `blacklist squashfs` in the `/etc/modprobe.d/` directory
- Unload `squashfs` from the kernel

If available in ANY installed kernel:

- Create a file with `blacklist squashfs` in the `/etc/modprobe.d/` directory

If the kernel module is not available on the system or pre-compiled into the kernel:

- No remediation is necessary

```
#!/usr/bin/env bash

{
  l_mname="squashfs" # set module name
  l_mtype="fs" # set module type
  l_mpath="/lib/modules/**/kernel/$l_mtype"
  l_mpname="$(tr '-' '_' <<< "$l_mname")"
  l_mndir="$(tr '-' '/' <<< "$l_mname")"

  module_loadable_fix()
  {
    # If the module is currently loadable, add "install {MODULE_NAME} /bin/false" to a file in
    "/etc/modprobe.d"
    l_loadable="$(modprobe -n -v "$l_mname")"
    [ "$(wc -l <<< "$l_loadable")" -gt "1" ] && l_loadable="$(grep -P --
    "(^\\h*install|\\b$l_mname)\\b" <<< "$l_loadable")"
    if ! grep -Pq -- '^\\h*install \\bin/(true|false)' <<< "$l_loadable"; then
      echo -e "\\n - setting module: \"$l_mname\" to be not loadable"
      echo -e "install $l_mname /bin/false" >> /etc/modprobe.d/"$l_mpname".conf
    fi
  }

  module_loaded_fix()
  {
    # If the module is currently loaded, unload the module
    if lsmod | grep "$l_mname" > /dev/null 2>&1; then
      echo -e "\\n - unloading module \"$l_mname\""
      modprobe -r "$l_mname"
    fi
  }

  module_deny_fix()
  {
    # If the module isn't deny listed, denylist the module
    if ! modprobe --showconfig | grep -Pq -- "^\\h*blacklist\\h+$l_mpname\\b"; then
      echo -e "\\n - deny listing \"$l_mname\""
      echo -e "blacklist $l_mname" >> /etc/modprobe.d/"$l_mpname".conf
    fi
  }

  # Check if the module exists on the system
  for l_mdir in $l_mpath; do
    if [ -d "$l_mdir/$l_mndir" ] && [ -n "$(ls -A $l_mdir/$l_mndir)" ]; then
      echo -e "\\n - module: \"$l_mname\" exists in \"$l_mdir\"\\n - checking if disabled..."
      module_deny_fix
      if [ "$l_mdir" = "/lib/modules/$(uname -r)/kernel/$l_mtype" ]; then
        module_loadable_fix
        module_loaded_fix
      fi
    else
      echo -e "\\n - module: \"$l_mname\" doesn't exist in \"$l_mdir\"\\n"
    fi
  done
  echo -e "\\n - remediation of module: \"$l_mname\" complete\\n"
}

```

References:

1. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1005, T1005.000	TA0005	M1050

1.1.1.2 Ensure mounting of udf filesystems is disabled (Automated)

Profile Applicability:

- Level 2 - Server

Description:

The `udf` filesystem type is the universal disk format used to implement ISO/IEC 13346 and ECMA-167 specifications. This is an open vendor filesystem type for data storage on a broad range of media. This filesystem type is necessary to support writing DVDs and newer optical disc formats.

Rationale:

Removing support for unneeded filesystem types reduces the local attack surface of the system. If this filesystem type is not needed, disable it.

Audit:

Run the following script to verify the `udf` module is disabled:

If the module is available in the running kernel:

- An entry including `/bin/true` or `/bin/false` exists in a file within the `/etc/modprobe.d/` directory
- The module is deny listed in a file within the `/etc/modprobe.d/` directory
- The module is not loaded in the kernel

If available in ANY installed kernel:

- The module is deny listed in a file within the `/etc/modprobe.d/` directory

If the kernel module is not available on the system, or pre-compiled into the kernel:

- No additional configuration is necessary

```

#!/usr/bin/env bash

{
    l_output="" l_output2="" l_output3="" l_dl="" # Unset output variables
    l_mname="udf" # set module name
    l_mtype="fs" # set module type
    l_searchloc="/lib/modprobe.d/*.conf /usr/local/lib/modprobe.d/*.conf /run/modprobe.d/*.conf
/etc/modprobe.d/*.conf"
    l_mpath="/lib/modules/**/kernel/$l_mtype"
    l_mpname="$(tr '-' '/' <<< "$l_mname")"
    l_mndir="$(tr '-' '/' <<< "$l_mname")"

    module_loadable_chk()
    {
        # Check if the module is currently loadable
        l_loadable="$(modprobe -n -v "$l_mname")"
        [ "$l_loadable" -gt "1" ] && l_loadable="$(grep -P --
"^(h*install|b$l_mname)b" <<< "$l_loadable")"
        if grep -Pq -- '^(h*install |bin|/true|false)' <<< "$l_loadable"; then
            l_output="$l_output\n - module: \"$l_mname\" is not loadable: \"$l_loadable\""
        else
            l_output2="$l_output2\n - module: \"$l_mname\" is loadable: \"$l_loadable\""
        fi
    }

    module_loaded_chk()
    {
        # Check if the module is currently loaded
        if ! lsmod | grep "$l_mname" > /dev/null 2>&1; then
            l_output="$l_output\n - module: \"$l_mname\" is not loaded"
        else
            l_output2="$l_output2\n - module: \"$l_mname\" is loaded"
        fi
    }

    module_deny_chk()
    {
        # Check if the module is deny listed
        l_dl="y"
        if modprobe --showconfig | grep -Pq -- '^(h*blacklist|h+"$l_mpname"'b'; then
            l_output="$l_output\n - module: \"$l_mname\" is deny listed in: \"$(grep -Pls --
"^(h*blacklist|h+$l_mname)b" $l_searchloc)\""
        else
            l_output2="$l_output2\n - module: \"$l_mname\" is not deny listed"
        fi
    }

    # Check if the module exists on the system
    for l_mdir in $l_mpath; do
        if [ -d "$l_mdir/$l_mndir" ] && [ -n "$(ls -A $l_mdir/$l_mndir)" ]; then
            l_output3="$l_output3\n - \"$l_mdir\""
            [ "$l_dl" != "y" ] && module_deny_chk
            if [ "$l_mdir" = "/lib/modules/$(uname -r)/kernel/$l_mtype" ]; then
                module_loadable_chk
                module_loaded_chk
            fi
        else
            l_output="$l_output\n - module: \"$l_mname\" doesn't exist in \"$l_mdir\""
        fi
    done

    # Report results. If no failures output in l_output2, we pass
    [ -n "$l_output3" ] && echo -e "\n\n -- INFO --\n - module: \"$l_mname\" exists in:$l_output3"
    if [ -z "$l_output2" ]; then
        echo -e "\n- Audit Result:\n  ** PASS **\n$l_output\n"
    else
        echo -e "\n- Audit Result:\n  ** FAIL **\n - Reason(s) for audit failure:\n$l_output2\n"
        [ -n "$l_output" ] && echo -e "\n- Correctly set:\n$l_output\n"
    fi
}

```

Remediation:

Run the following script to disable the `udf` module:

If the module is available in the running kernel:

- Create a file with `install udf /bin/false` in the `/etc/modprobe.d/` directory
- Create a file with `blacklist udf` in the `/etc/modprobe.d/` directory
- Unload `udf` from the kernel

If available in ANY installed kernel:

- Create a file with `blacklist udf` in the `/etc/modprobe.d/` directory

If the kernel module is not available on the system or pre-compiled into the kernel:

- No remediation is necessary

```
#!/usr/bin/env bash

{
  l_mname="udf" # set module name
  l_mtype="fs" # set module type
  l_mpath="/lib/modules/**/kernel/$l_mtype"
  l_mpname="$(tr '-' '/' <<< "$l_mname")"
  l_mndir="$(tr '-' '/' <<< "$l_mname")"

  module_loadable_fix()
  {
    # If the module is currently loadable, add "install {MODULE_NAME} /bin/false" to a file in
    "/etc/modprobe.d"
    l_loadable="$(modprobe -n -v "$l_mname")"
    [ "$(wc -l <<< "$l_loadable")" -gt "1" ] && l_loadable="$(grep -P --
    "(^\\h*install|\\b$l_mname\\b)" <<< "$l_loadable")"
    if ! grep -Pq -- '^\\h*install \\bin/(true|false)' <<< "$l_loadable"; then
      echo -e "\\n - setting module: \"$l_mname\" to be not loadable"
      echo -e "install $l_mname /bin/false" >> /etc/modprobe.d/"$l_mpname".conf
    fi
  }
  module_loaded_fix()
  {
    # If the module is currently loaded, unload the module
    if lsmod | grep "$l_mname" > /dev/null 2>&1; then
      echo -e "\\n - unloading module \"$l_mname\""
      modprobe -r "$l_mname"
    fi
  }
  module_deny_fix()
  {
    # If the module isn't deny listed, denylist the module
    if ! modprobe --showconfig | grep -Pq -- "^\\h*blacklist\\h+$l_mpname\\b"; then
      echo -e "\\n - deny listing \"$l_mname\""
      echo -e "blacklist $l_mname" >> /etc/modprobe.d/"$l_mpname".conf
    fi
  }
  # Check if the module exists on the system
  for l_mdir in $l_mpath; do
    if [ -d "$l_mdir/$l_mndir" ] && [ -n "$(ls -A $l_mdir/$l_mndir)" ]; then
      echo -e "\\n - module: \"$l_mname\" exists in \"$l_mdir\"\\n - checking if disabled..."
      module_deny_fix
      if [ "$l_mdir" = "/lib/modules/$(uname -r)/kernel/$l_mtype" ]; then
        module_loadable_fix
        module_loaded_fix
      fi
    else
      echo -e "\\n - module: \"$l_mname\" doesn't exist in \"$l_mdir\"\\n"
    fi
  done
  echo -e "\\n - remediation of module: \"$l_mname\" complete\\n"
}

```

References:

1. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1005, T1005.000	TA0005	M1050

1.1.1.3 Ensure mounting of cramfs filesystems is disabled (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `cramfs` filesystem type is a compressed read-only Linux filesystem embedded in small footprint systems. A `cramfs` image can be used without having to first decompress the image.

Rationale:

Removing support for unneeded filesystem types reduces the local attack surface of the system. If this filesystem type is not needed, disable it.

Audit:

Run the following script to verify the `cramfs` module is disabled:

If the module is available in the running kernel:

- An entry including `/bin/true` or `/bin/false` exists in a file within the `/etc/modprobe.d/` directory
- The module is deny listed in a file within the `/etc/modprobe.d/` directory
- The module is not loaded in the kernel

If available in ANY installed kernel:

- The module is deny listed in a file within the `/etc/modprobe.d/` directory

If the kernel module is not available on the system, or pre-compiled into the kernel:

- No additional configuration is necessary

```

#!/usr/bin/env bash

{
    l_output="" l_output2="" l_output3="" l_dl="" # Unset output variables
    l_mname="cramfs" # set module name
    l_mtype="fs" # set module type
    l_searchloc="/lib/modprobe.d/*.conf /usr/local/lib/modprobe.d/*.conf /run/modprobe.d/*.conf
/etc/modprobe.d/*.conf"
    l_mpath="/lib/modules/**/kernel/$l_mtype"
    l_mpname="$(tr '-' '/' <<< "$l_mname")"
    l_mndir="$(tr '-' '/' <<< "$l_mname")"

    module_loadable_chk()
    {
        # Check if the module is currently loadable
        l_loadable="$(modprobe -n -v "$l_mname")"
        [ "$l_loadable" -gt "1" ] && l_loadable="$(grep -P --
"^(^\\h*install|\\b$l_mname)\\b" <<< "$l_loadable")"
        if grep -Pq -- '^\\h*install \\bin\\(true|false\\)' <<< "$l_loadable"; then
            l_output="$l_output\\n - module: \"$l_mname\" is not loadable: \"$l_loadable\\n\"
        else
            l_output2="$l_output2\\n - module: \"$l_mname\" is loadable: \"$l_loadable\\n\"
        fi
    }

    module_loaded_chk()
    {
        # Check if the module is currently loaded
        if ! lsmod | grep "$l_mname" > /dev/null 2>&1; then
            l_output="$l_output\\n - module: \"$l_mname\" is not loaded"
        else
            l_output2="$l_output2\\n - module: \"$l_mname\" is loaded"
        fi
    }

    module_deny_chk()
    {
        # Check if the module is deny listed
        l_dl="y"
        if modprobe --showconfig | grep -Pq -- '^\\h*blacklist\\h+\"$l_mpname\"\\b'; then
            l_output="$l_output\\n - module: \"$l_mname\" is deny listed in: \"$(grep -Pls --
"^(\\h*blacklist\\h+$l_mname\\b) $l_searchloc)\\n\"
        else
            l_output2="$l_output2\\n - module: \"$l_mname\" is not deny listed"
        fi
    }

    # Check if the module exists on the system
    for l_mdir in $l_mpath; do
        if [ -d "$l_mdir/$l_mndir" ] && [ -n "$(ls -A $l_mdir/$l_mndir)" ]; then
            l_output3="$l_output3\\n - \"$l_mdir\\n\"
            [ "$l_dl" != "y" ] && module_deny_chk
            if [ "$l_mdir" = "/lib/modules/$(uname -r)/kernel/$l_mtype" ]; then
                module_loadable_chk
                module_loaded_chk
            fi
        else
            l_output="$l_output\\n - module: \"$l_mname\" doesn't exist in \"$l_mdir\\n\"
        fi
    done

    # Report results. If no failures output in l_output2, we pass
    [ -n "$l_output3" ] && echo -e "\\n\\n -- INFO --\\n - module: \"$l_mname\" exists in:$l_output3"
    if [ -z "$l_output2" ]; then
        echo -e "\\n- Audit Result:\\n ** PASS **\\n$l_output\\n"
    else
        echo -e "\\n- Audit Result:\\n ** FAIL **\\n - Reason(s) for audit failure:\\n$l_output2\\n"
        [ -n "$l_output" ] && echo -e "\\n- Correctly set:\\n$l_output\\n"
    fi
}

```

Remediation:

Run the following script to disable the `cramfs` module:

If the module is available in the running kernel:

- Create a file with `install cramfs /bin/false` in the `/etc/modprobe.d/` directory
- Create a file with `blacklist cramfs` in the `/etc/modprobe.d/` directory
- Unload `cramfs` from the kernel

If available in ANY installed kernel:

- Create a file with `blacklist cramfs` in the `/etc/modprobe.d/` directory

If the kernel module is not available on the system or pre-compiled into the kernel:

- No remediation is necessary


```
#!/usr/bin/env bash

{
  l_mname="cramfs" # set module name
  l_mtype="fs" # set module type
  l_mpath="/lib/modules/**/kernel/$l_mtype"
  l_mpname="$(tr '-' '_' <<< "$l_mname")"
  l_mndir="$(tr '-' '/' <<< "$l_mname")"

  module_loadable_fix()
  {
    # If the module is currently loadable, add "install {MODULE_NAME} /bin/false" to a file in
    "/etc/modprobe.d"
    l_loadable="$(modprobe -n -v "$l_mname")"
    [ "$(wc -l <<< "$l_loadable")" -gt "1" ] && l_loadable="$(grep -P --
    "(^\\h*install|\\b$l_mname)\\b" <<< "$l_loadable")"
    if ! grep -Pq -- '^\\h*install \\bin/(true|false)' <<< "$l_loadable"; then
      echo -e "\\n - setting module: \"$l_mname\" to be not loadable"
      echo -e "install $l_mname /bin/false" >> /etc/modprobe.d/"$l_mpname".conf
    fi
  }

  module_loaded_fix()
  {
    # If the module is currently loaded, unload the module
    if lsmod | grep "$l_mname" > /dev/null 2>&1; then
      echo -e "\\n - unloading module \"$l_mname\""
      modprobe -r "$l_mname"
    fi
  }

  module_deny_fix()
  {
    # If the module isn't deny listed, denylist the module
    if ! modprobe --showconfig | grep -Pq -- "^\\h*blacklist\\h+$l_mpname\\b"; then
      echo -e "\\n - deny listing \"$l_mname\""
      echo -e "blacklist $l_mname" >> /etc/modprobe.d/"$l_mpname".conf
    fi
  }

  # Check if the module exists on the system
  for l_mdir in $l_mpath; do
    if [ -d "$l_mdir/$l_mndir" ] && [ -n "$(ls -A $l_mdir/$l_mndir)" ]; then
      echo -e "\\n - module: \"$l_mname\" exists in \"$l_mdir\"\\n - checking if disabled..."
      module_deny_fix
      if [ "$l_mdir" = "/lib/modules/$(uname -r)/kernel/$l_mtype" ]; then
        module_loadable_fix
        module_loaded_fix
      fi
    else
      echo -e "\\n - module: \"$l_mname\" doesn't exist in \"$l_mdir\"\\n"
    fi
  done
  echo -e "\\n - remediation of module: \"$l_mname\" complete\\n"
}

```

References:

1. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1005, T1005.000	TA0005	M1050

1.1.1.4 Ensure mounting of freevxfs filesystems is disabled (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `freevxfs` filesystem type is a free version of the Veritas type filesystem. This is the primary filesystem type for HP-UX operating systems.

Rationale:

Removing support for unneeded filesystem types reduces the local attack surface of the system. If this filesystem type is not needed, disable it.

Audit:

Run the following script to verify the `freevxfs` module is disabled:
If the module is available in the running kernel:

- An entry including `/bin/true` or `/bin/false` exists in a file within the `/etc/modprobe.d/` directory
- The module is deny listed in a file within the `/etc/modprobe.d/` directory
- The module is not loaded in the kernel

If available in ANY installed kernel:

- The module is deny listed in a file within the `/etc/modprobe.d/` directory

If the kernel module is not available on the system, or pre-compiled into the kernel:

- No additional configuration is necessary

```

#!/usr/bin/env bash

{
    l_output="" l_output2="" l_output3="" l_dl="" # Unset output variables
    l_mname="freevxfs" # set module name
    l_mtype="fs" # set module type
    l_searchloc="/lib/modprobe.d/*.conf /usr/local/lib/modprobe.d/*.conf /run/modprobe.d/*.conf
/etc/modprobe.d/*.conf"
    l_mpath="/lib/modules/**/kernel/$l_mtype"
    l_mpname="$(tr '-' '_' <<< "$l_mname")"
    l_mndir="$(tr '-' '/' <<< "$l_mname")"
    module_loadable_chk()
    {
        # Check if the module is currently loadable
        l_loadable="$(modprobe -n -v "$l_mname")"
        [ "$(wc -l <<< "$l_loadable")" -gt "1" ] && l_loadable="$(grep -P --
"^(^\\h*install|\\b$l_mname)\\b" <<< "$l_loadable")"
        if grep -Pq -- '^\\h*install \\bin/(true|false)' <<< "$l_loadable"; then
            l_output="$l_output\\n - module: \"$l_mname\" is not loadable: \"$l_loadable\\n"
        else
            l_output2="$l_output2\\n - module: \"$l_mname\" is loadable: \"$l_loadable\\n"
        fi
    }
    module_loaded_chk()
    {
        # Check if the module is currently loaded
        if ! lsmod | grep "$l_mname" > /dev/null 2>&1; then
            l_output="$l_output\\n - module: \"$l_mname\" is not loaded"
        else
            l_output2="$l_output2\\n - module: \"$l_mname\" is loaded"
        fi
    }
    module_deny_chk()
    {
        # Check if the module is deny listed
        l_dl="y"
        if modprobe --showconfig | grep -Pq -- '^\\h*blacklist\\h+' "$l_mpname" '\\b'; then
            l_output="$l_output\\n - module: \"$l_mname\" is deny listed in: \"$(grep -Pls --
"^(^\\h*blacklist\\h+$l_mname\\b" $l_searchloc)\\n"
        else
            l_output2="$l_output2\\n - module: \"$l_mname\" is not deny listed"
        fi
    }
    # Check if the module exists on the system
    for l_mdir in $l_mpath; do
        if [ -d "$l_mdir/$l_mndir" ] && [ -n "$(ls -A $l_mdir/$l_mndir)" ]; then
            l_output3="$l_output3\\n - \"$l_mdir\\n"
            [ "$l_dl" != "y" ] && module_deny_chk
            if [ "$l_mdir" = "/lib/modules/$(uname -r)/kernel/$l_mtype" ]; then
                module_loadable_chk
                module_loaded_chk
            fi
        else
            l_output="$l_output\\n - module: \"$l_mname\" doesn't exist in \"$l_mdir\\n"
        fi
    done
    # Report results. If no failures output in l_output2, we pass
    [ -n "$l_output3" ] && echo -e "\\n\\n -- INFO --\\n - module: \"$l_mname\" exists in:$l_output3"
    if [ -z "$l_output2" ]; then
        echo -e "\\n- Audit Result:\\n ** PASS **\\n$l_output\\n"
    else
        echo -e "\\n- Audit Result:\\n ** FAIL **\\n - Reason(s) for audit failure:\\n$l_output2\\n"
        [ -n "$l_output" ] && echo -e "\\n- Correctly set:\\n$l_output\\n"
    fi
}

```

Remediation:

Run the following script to disable the `freevxfs` module:

If the module is available in the running kernel:

- Create a file with `install freevxfs /bin/false` in the `/etc/modprobe.d/` directory
- Create a file with `blacklist freevxfs` in the `/etc/modprobe.d/` directory
- Unload `freevxfs` from the kernel

If available in ANY installed kernel:

- Create a file with `blacklist freevxfs` in the `/etc/modprobe.d/` directory

If the kernel module is not available on the system or pre-compiled into the kernel:

- No remediation is necessary

```
#!/usr/bin/env bash

{
  l_mname="freevxfs" # set module name
  l_mtype="fs" # set module type
  l_mpath="/lib/modules/**/kernel/$l_mtype"
  l_mpname="$(tr '-' '_' <<< "$l_mname")"
  l_mndir="$(tr '-' '/' <<< "$l_mname")"

  module_loadable_fix()
  {
    # If the module is currently loadable, add "install {MODULE_NAME} /bin/false" to a file in
    "/etc/modprobe.d"
    l_loadable="$(modprobe -n -v "$l_mname")"
    [ "$(wc -l <<< "$l_loadable")" -gt "1" ] && l_loadable="$(grep -P --
    "(^\\h*install|\\b$l_mname)\\b" <<< "$l_loadable")"
    if ! grep -Pq -- '^\\h*install \\bin/(true|false)' <<< "$l_loadable"; then
      echo -e "\\n - setting module: \"$l_mname\" to be not loadable"
      echo -e "install $l_mname /bin/false" >> /etc/modprobe.d/"$l_mpname".conf
    fi
  }
  module_loaded_fix()
  {
    # If the module is currently loaded, unload the module
    if lsmod | grep "$l_mname" > /dev/null 2>&1; then
      echo -e "\\n - unloading module \"$l_mname\""
      modprobe -r "$l_mname"
    fi
  }
  module_deny_fix()
  {
    # If the module isn't deny listed, denylist the module
    if ! modprobe --showconfig | grep -Pq -- "^\\h*blacklist\\h+$l_mpname\\b"; then
      echo -e "\\n - deny listing \"$l_mname\""
      echo -e "blacklist $l_mname" >> /etc/modprobe.d/"$l_mpname".conf
    fi
  }
  # Check if the module exists on the system
  for l_mdir in $l_mpath; do
    if [ -d "$l_mdir/$l_mndir" ] && [ -n "$(ls -A $l_mdir/$l_mndir)" ]; then
      echo -e "\\n - module: \"$l_mname\" exists in \"$l_mdir\"\\n - checking if disabled..."
      module_deny_fix
      if [ "$l_mdir" = "/lib/modules/$(uname -r)/kernel/$l_mtype" ]; then
        module_loadable_fix
        module_loaded_fix
      fi
    else
      echo -e "\\n - module: \"$l_mname\" doesn't exist in \"$l_mdir\"\\n"
    fi
  done
  echo -e "\\n - remediation of module: \"$l_mname\" complete\\n"
}

```

References:

1. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	9.2 <u>Use DNS Filtering Services</u> Use DNS filtering services on all enterprise assets to block access to known malicious domains.			
v7	4.8 <u>Log and Alert on Changes to Administrative Group Membership</u> Configure systems to issue a log entry and alert when an account is added to or removed from any group assigned administrative privileges.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1005, T1005.000	TA0005	M1050

1.1.1.5 Ensure mounting of jffs2 filesystems is disabled (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `jffs2` (journaling flash filesystem 2) filesystem type is a log-structured filesystem used in flash memory devices.

Rationale:

Removing support for unneeded filesystem types reduces the local attack surface of the system. If this filesystem type is not needed, disable it.

Audit:

Run the following script to verify the `jffs2` module is disabled:
If the module is available in the running kernel:

- An entry including `/bin/true` or `/bin/false` exists in a file within the `/etc/modprobe.d/` directory
- The module is deny listed in a file within the `/etc/modprobe.d/` directory
- The module is not loaded in the kernel

If available in ANY installed kernel:

- The module is deny listed in a file within the `/etc/modprobe.d/` directory

If the kernel module is not available on the system, or pre-compiled into the kernel:

- No additional configuration is necessary


```

#!/usr/bin/env bash

{
    l_output="" l_output2="" l_output3="" l_dl="" # Unset output variables
    l_mname="jffs2" # set module name
    l_mtype="fs" # set module type
    l_searchloc="/lib/modprobe.d/*.conf /usr/local/lib/modprobe.d/*.conf /run/modprobe.d/*.conf
/etc/modprobe.d/*.conf"
    l_mpath="/lib/modules/**/kernel/$l_mtype"
    l_mpname="$(tr '-' '/' <<< "$l_mname")"
    l_mndir="$(tr '-' '/' <<< "$l_mname")"

    module_loadable_chk()
    {
        # Check if the module is currently loadable
        l_loadable="$(modprobe -n -v "$l_mname")"
        [ "$l_loadable" -gt "1" ] && l_loadable="$(grep -P --
"^(^\\h*install|\\b$l_mname)\\b" <<< "$l_loadable")"
        if grep -Pq -- '^\\h*install \\bin\\(true|false\\)' <<< "$l_loadable"; then
            l_output="$l_output\\n - module: \"$l_mname\" is not loadable: \"$l_loadable\\n\"
        else
            l_output2="$l_output2\\n - module: \"$l_mname\" is loadable: \"$l_loadable\\n\"
        fi
    }

    module_loaded_chk()
    {
        # Check if the module is currently loaded
        if ! lsmod | grep "$l_mname" > /dev/null 2>&1; then
            l_output="$l_output\\n - module: \"$l_mname\" is not loaded"
        else
            l_output2="$l_output2\\n - module: \"$l_mname\" is loaded"
        fi
    }

    module_deny_chk()
    {
        # Check if the module is deny listed
        l_dl="y"
        if modprobe --showconfig | grep -Pq -- '^\\h*blacklist\\h+'$l_mpname'"\\b'; then
            l_output="$l_output\\n - module: \"$l_mname\" is deny listed in: \"$(grep -Pls --
"^(\\h*blacklist\\h+$l_mname\\b" $l_searchloc)\\n\"
        else
            l_output2="$l_output2\\n - module: \"$l_mname\" is not deny listed"
        fi
    }

    # Check if the module exists on the system
    for l_mdir in $l_mpath; do
        if [ -d "$l_mdir/$l_mndir" ] && [ -n "$(ls -A $l_mdir/$l_mndir)" ]; then
            l_output3="$l_output3\\n - \"$l_mdir\\n\"
            [ "$l_dl" != "y" ] && module_deny_chk
            if [ "$l_mdir" = "/lib/modules/$(uname -r)/kernel/$l_mtype" ]; then
                module_loadable_chk
                module_loaded_chk
            fi
        else
            l_output="$l_output\\n - module: \"$l_mname\" doesn't exist in \"$l_mdir\\n\"
        fi
    done

    # Report results. If no failures output in l_output2, we pass
    [ -n "$l_output3" ] && echo -e "\\n\\n -- INFO --\\n - module: \"$l_mname\" exists in:$l_output3"
    if [ -z "$l_output2" ]; then
        echo -e "\\n- Audit Result:\\n ** PASS **\\n$l_output\\n"
    else
        echo -e "\\n- Audit Result:\\n ** FAIL **\\n - Reason(s) for audit failure:\\n$l_output2\\n"
        [ -n "$l_output" ] && echo -e "\\n- Correctly set:\\n$l_output\\n"
    fi
}

```

Remediation:

Run the following script to disable the `jffs2` module:

If the module is available in the running kernel:

- Create a file with `install jffs2 /bin/false` in the `/etc/modprobe.d/` directory
- Create a file with `blacklist jffs2` in the `/etc/modprobe.d/` directory
- Unload `jffs2` from the kernel

If available in ANY installed kernel:

- Create a file with `blacklist jffs2` in the `/etc/modprobe.d/` directory

If the kernel module is not available on the system or pre-compiled into the kernel:

- No remediation is necessary

```
#!/usr/bin/env bash

{
  l_mname="jffs2" # set module name
  l_mtype="fs" # set module type
  l_mpath="/lib/modules/**/kernel/$l_mtype"
  l_mpname="$(tr '-' '_' <<< "$l_mname")"
  l_mndir="$(tr '-' '/' <<< "$l_mname")"

  module_loadable_fix()
  {
    # If the module is currently loadable, add "install {MODULE_NAME} /bin/false" to a file in
    "/etc/modprobe.d"
    l_loadable="$(modprobe -n -v "$l_mname")"
    [ "$(wc -l <<< "$l_loadable")" -gt "1" ] && l_loadable="$(grep -P --
    "(^\\h*install|\\b$l_mname\\b)" <<< "$l_loadable")"
    if ! grep -Pq -- '^\\h*install \\bin/(true|false)' <<< "$l_loadable"; then
      echo -e "\\n - setting module: \"$l_mname\" to be not loadable"
      echo -e "install $l_mname /bin/false" >> /etc/modprobe.d/"$l_mpname".conf
    fi
  }
  module_loaded_fix()
  {
    # If the module is currently loaded, unload the module
    if lsmod | grep "$l_mname" > /dev/null 2>&1; then
      echo -e "\\n - unloading module \"$l_mname\""
      modprobe -r "$l_mname"
    fi
  }
  module_deny_fix()
  {
    # If the module isn't deny listed, denylist the module
    if ! modprobe --showconfig | grep -Pq -- "^\\h*blacklist\\h+$l_mpname\\b"; then
      echo -e "\\n - deny listing \"$l_mname\""
      echo -e "blacklist $l_mname" >> /etc/modprobe.d/"$l_mpname".conf
    fi
  }
  # Check if the module exists on the system
  for l_mdir in $l_mpath; do
    if [ -d "$l_mdir/$l_mndir" ] && [ -n "$(ls -A $l_mdir/$l_mndir)" ]; then
      echo -e "\\n - module: \"$l_mname\" exists in \"$l_mdir\"\\n - checking if disabled..."
      module_deny_fix
      if [ "$l_mdir" = "/lib/modules/$(uname -r)/kernel/$l_mtype" ]; then
        module_loadable_fix
        module_loaded_fix
      fi
    else
      echo -e "\\n - module: \"$l_mname\" doesn't exist in \"$l_mdir\"\\n"
    fi
  done
  echo -e "\\n - remediation of module: \"$l_mname\" complete\\n"
}

```

References:

1. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1005, T1005.000	TA0005	M1050

1.1.1.6 Ensure mounting of hfs filesystems is disabled (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `hfs` filesystem type is a hierarchical filesystem that allows you to mount Mac OS filesystems.

Rationale:

Removing support for unneeded filesystem types reduces the local attack surface of the system. If this filesystem type is not needed, disable it.

Audit:

Run the following script to verify the `hfs` module is disabled:

If the module is available in the running kernel:

- An entry including `/bin/true` or `/bin/false` exists in a file within the `/etc/modprobe.d/` directory
- The module is deny listed in a file within the `/etc/modprobe.d/` directory
- The module is not loaded in the kernel

If available in ANY installed kernel:

- The module is deny listed in a file within the `/etc/modprobe.d/` directory

If the kernel module is not available on the system, or pre-compiled into the kernel:

- No additional configuration is necessary

```

#!/usr/bin/env bash

{
    l_output="" l_output2="" l_output3="" l_dl="" # Unset output variables
    l_mname="hfs" # set module name
    l_mtype="fs" # set module type
    l_searchloc="/lib/modprobe.d/*.conf /usr/local/lib/modprobe.d/*.conf /run/modprobe.d/*.conf
/etc/modprobe.d/*.conf"
    l_mpath="/lib/modules/**/kernel/$l_mtype"
    l_mpname="$(tr '-' '/' <<< "$l_mname")"
    l_mndir="$(tr '-' '/' <<< "$l_mname")"

    module_loadable_chk()
    {
        # Check if the module is currently loadable
        l_loadable="$(modprobe -n -v "$l_mname")"
        [ "$l_loadable" -gt "1" ] && l_loadable="$(grep -P --
"^(h*install|b$l_mname)b" <<< "$l_loadable")"
        if grep -Pq -- '^(h*install |bin|/true|false)' <<< "$l_loadable"; then
            l_output="$l_output\n - module: \"$l_mname\" is not loadable: \"$l_loadable\""
        else
            l_output2="$l_output2\n - module: \"$l_mname\" is loadable: \"$l_loadable\""
        fi
    }

    module_loaded_chk()
    {
        # Check if the module is currently loaded
        if ! lsmod | grep "$l_mname" > /dev/null 2>&1; then
            l_output="$l_output\n - module: \"$l_mname\" is not loaded"
        else
            l_output2="$l_output2\n - module: \"$l_mname\" is loaded"
        fi
    }

    module_deny_chk()
    {
        # Check if the module is deny listed
        l_dl="y"
        if modprobe --showconfig | grep -Pq -- '^(h*blacklist|h+"$l_mpname"'b'; then
            l_output="$l_output\n - module: \"$l_mname\" is deny listed in: \"$(grep -Pls --
"^(h*blacklist|h+$l_mname)b" $l_searchloc)\""
        else
            l_output2="$l_output2\n - module: \"$l_mname\" is not deny listed"
        fi
    }

    # Check if the module exists on the system
    for l_mdir in $l_mpath; do
        if [ -d "$l_mdir/$l_mndir" ] && [ -n "$(ls -A $l_mdir/$l_mndir)" ]; then
            l_output3="$l_output3\n - \"$l_mdir\""
            [ "$l_dl" != "y" ] && module_deny_chk
            if [ "$l_mdir" = "/lib/modules/$(uname -r)/kernel/$l_mtype" ]; then
                module_loadable_chk
                module_loaded_chk
            fi
        else
            l_output="$l_output\n - module: \"$l_mname\" doesn't exist in \"$l_mdir\""
        fi
    done

    # Report results. If no failures output in l_output2, we pass
    [ -n "$l_output3" ] && echo -e "\n\n -- INFO --\n - module: \"$l_mname\" exists in:$l_output3"
    if [ -z "$l_output2" ]; then
        echo -e "\n- Audit Result:\n  ** PASS **\n$l_output\n"
    else
        echo -e "\n- Audit Result:\n  ** FAIL **\n - Reason(s) for audit failure:\n$l_output2\n"
        [ -n "$l_output" ] && echo -e "\n- Correctly set:\n$l_output\n"
    fi
}

```

Remediation:

Run the following script to disable the `hfs` module:

If the module is available in the running kernel:

- Create a file with `install hfs /bin/false` in the `/etc/modprobe.d/` directory
- Create a file with `blacklist hfs` in the `/etc/modprobe.d/` directory
- Unload `hfs` from the kernel

If available in ANY installed kernel:

- Create a file with `blacklist hfs` in the `/etc/modprobe.d/` directory

If the kernel module is not available on the system or pre-compiled into the kernel:

- No remediation is necessary

```
#!/usr/bin/env bash

{
    l_mname="hfs" # set module name
    l_mtype="fs" # set module type
    l_mpath="/lib/modules/**/kernel/$l_mtype"
    l_mpname="$(tr '-' '/' <<< "$l_mname")"
    l_mndir="$(tr '-' '/' <<< "$l_mname")"

    module_loadable_fix()
    {
        # If the module is currently loadable, add "install {MODULE_NAME} /bin/false" to a file in
        "/etc/modprobe.d"
        l_loadable="$(modprobe -n -v "$l_mname")"
        [ "$(wc -l <<< "$l_loadable")" -gt "1" ] && l_loadable="$(grep -P --
        "(^\\h*install|\\b$l_mname\\b)" <<< "$l_loadable")"
        if ! grep -Pq -- '^\\h*install \\bin/(true|false)' <<< "$l_loadable"; then
            echo -e "\\n - setting module: \"$l_mname\" to be not loadable"
            echo -e "install $l_mname /bin/false" >> /etc/modprobe.d/"$l_mpname".conf
        fi
    }

    module_loaded_fix()
    {
        # If the module is currently loaded, unload the module
        if lsmod | grep "$l_mname" > /dev/null 2>&1; then
            echo -e "\\n - unloading module \"$l_mname\""
            modprobe -r "$l_mname"
        fi
    }

    module_deny_fix()
    {
        # If the module isn't deny listed, denylist the module
        if ! modprobe --showconfig | grep -Pq -- "^\\h*blacklist\\h+$l_mpname\\b"; then
            echo -e "\\n - deny listing \"$l_mname\""
            echo -e "blacklist $l_mname" >> /etc/modprobe.d/"$l_mpname".conf
        fi
    }

    # Check if the module exists on the system
    for l_mdir in $l_mpath; do
        if [ -d "$l_mdir/$l_mndir" ] && [ -n "$(ls -A $l_mdir/$l_mndir)" ]; then
            echo -e "\\n - module: \"$l_mname\" exists in \"$l_mdir\"\\n - checking if disabled..."
            module_deny_fix
            if [ "$l_mdir" = "/lib/modules/$(uname -r)/kernel/$l_mtype" ]; then
                module_loadable_fix
                module_loaded_fix
            fi
        else
            echo -e "\\n - module: \"$l_mname\" doesn't exist in \"$l_mdir\"\\n"
        fi
    done
    echo -e "\\n - remediation of module: \"$l_mname\" complete\\n"
}

```

References:

1. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1005, T1005.000	TA0005	M1050

1.1.1.7 Ensure mounting of hfsplus filesystems is disabled (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `hfsplus` filesystem type is a hierarchical filesystem designed to replace `hfs` that allows you to mount Mac OS filesystems.

Rationale:

Removing support for unneeded filesystem types reduces the local attack surface of the system. If this filesystem type is not needed, disable it.

Audit:

Run the following script to verify the `hfsplus` module is disabled:

If the module is available in the running kernel:

- An entry including `/bin/true` or `/bin/false` exists in a file within the `/etc/modprobe.d/` directory
- The module is deny listed in a file within the `/etc/modprobe.d/` directory
- The module is not loaded in the kernel

If available in ANY installed kernel:

- The module is deny listed in a file within the `/etc/modprobe.d/` directory

If the kernel module is not available on the system, or pre-compiled into the kernel:

- No additional configuration is necessary

```

#!/usr/bin/env bash

{
    l_output="" l_output2="" l_output3="" l_dl="" # Unset output variables
    l_mname="hfsplus" # set module name
    l_mtype="fs" # set module type
    l_searchloc="/lib/modprobe.d/*.conf /usr/local/lib/modprobe.d/*.conf /run/modprobe.d/*.conf
/etc/modprobe.d/*.conf"
    l_mpath="/lib/modules/**/kernel/$l_mtype"
    l_mpname="$(tr '-' '/' <<< "$l_mname")"
    l_mndir="$(tr '-' '/' <<< "$l_mname")"

    module_loadable_chk()
    {
        # Check if the module is currently loadable
        l_loadable="$(modprobe -n -v "$l_mname")"
        [ "$l_loadable" -gt "1" ] && l_loadable="$(grep -P --
"^(^\\h*install|\\b$l_mname)\\b" <<< "$l_loadable")"
        if grep -Pq -- '^\\h*install \\bin\\(true|false\\)' <<< "$l_loadable"; then
            l_output="$l_output\\n - module: \"$l_mname\" is not loadable: \"$l_loadable\\n\"
        else
            l_output2="$l_output2\\n - module: \"$l_mname\" is loadable: \"$l_loadable\\n\"
        fi
    }

    module_loaded_chk()
    {
        # Check if the module is currently loaded
        if ! lsmod | grep "$l_mname" > /dev/null 2>&1; then
            l_output="$l_output\\n - module: \"$l_mname\" is not loaded"
        else
            l_output2="$l_output2\\n - module: \"$l_mname\" is loaded"
        fi
    }

    module_deny_chk()
    {
        # Check if the module is deny listed
        l_dl="y"
        if modprobe --showconfig | grep -Pq -- '^\\h*blacklist\\h+'$l_mpname'"\\b'; then
            l_output="$l_output\\n - module: \"$l_mname\" is deny listed in: \"$(grep -Pls --
"^(\\h*blacklist\\h+$l_mname\\b" $l_searchloc)\\n\"
        else
            l_output2="$l_output2\\n - module: \"$l_mname\" is not deny listed"
        fi
    }

    # Check if the module exists on the system
    for l_mdir in $l_mpath; do
        if [ -d "$l_mdir/$l_mndir" ] && [ -n "$(ls -A $l_mdir/$l_mndir)" ]; then
            l_output3="$l_output3\\n - \"$l_mdir\\n\"
            [ "$l_dl" != "y" ] && module_deny_chk
            if [ "$l_mdir" = "/lib/modules/$(uname -r)/kernel/$l_mtype" ]; then
                module_loadable_chk
                module_loaded_chk
            fi
        else
            l_output="$l_output\\n - module: \"$l_mname\" doesn't exist in \"$l_mdir\\n\"
        fi
    done

    # Report results. If no failures output in l_output2, we pass
    [ -n "$l_output3" ] && echo -e "\\n\\n -- INFO --\\n - module: \"$l_mname\" exists in:$l_output3"
    if [ -z "$l_output2" ]; then
        echo -e "\\n- Audit Result:\\n ** PASS **\\n$l_output\\n"
    else
        echo -e "\\n- Audit Result:\\n ** FAIL **\\n - Reason(s) for audit failure:\\n$l_output2\\n"
        [ -n "$l_output" ] && echo -e "\\n- Correctly set:\\n$l_output\\n"
    fi
}

```

Remediation:

Run the following script to disable the `hfsplus` module:

If the module is available in the running kernel:

- Create a file with `install hfsplus /bin/false` in the `/etc/modprobe.d/` directory
- Create a file with `blacklist hfsplus` in the `/etc/modprobe.d/` directory
- Unload `hfsplus` from the kernel

If available in ANY installed kernel:

- Create a file with `blacklist hfsplus` in the `/etc/modprobe.d/` directory

If the kernel module is not available on the system or pre-compiled into the kernel:

- No remediation is necessary

```
#!/usr/bin/env bash

{
    l_mname="hfsplus" # set module name
    l_mtype="fs" # set module type
    l_mpath="/lib/modules/**/kernel/$l_mtype"
    l_mpname="$(tr '-' '_' <<< "$l_mname")"
    l_mndir="$(tr '-' '/' <<< "$l_mname")"

    module_loadable_fix()
    {
        # If the module is currently loadable, add "install {MODULE_NAME} /bin/false" to a file in
        "/etc/modprobe.d"
        l_loadable="$(modprobe -n -v "$l_mname")"
        [ "$(wc -l <<< "$l_loadable")" -gt "1" ] && l_loadable="$(grep -P --
        "(^\\h*install|\\b$l_mname\\b)" <<< "$l_loadable")"
        if ! grep -Pq -- '^\\h*install \\bin/(true|false)' <<< "$l_loadable"; then
            echo -e "\\n - setting module: \"$l_mname\" to be not loadable"
            echo -e "install $l_mname /bin/false" >> /etc/modprobe.d/"$l_mpname".conf
        fi
    }
    module_loaded_fix()
    {
        # If the module is currently loaded, unload the module
        if lsmod | grep "$l_mname" > /dev/null 2>&1; then
            echo -e "\\n - unloading module \"$l_mname\""
            modprobe -r "$l_mname"
        fi
    }
    module_deny_fix()
    {
        # If the module isn't deny listed, denylist the module
        if ! modprobe --showconfig | grep -Pq -- "^\\h*blacklist\\h+$l_mpname\\b"; then
            echo -e "\\n - deny listing \"$l_mname\""
            echo -e "blacklist $l_mname" >> /etc/modprobe.d/"$l_mpname".conf
        fi
    }
    # Check if the module exists on the system
    for l_mdir in $l_mpath; do
        if [ -d "$l_mdir/$l_mndir" ] && [ -n "$(ls -A $l_mdir/$l_mndir)" ]; then
            echo -e "\\n - module: \"$l_mname\" exists in \"$l_mdir\"\\n - checking if disabled..."
            module_deny_fix
            if [ "$l_mdir" = "/lib/modules/$(uname -r)/kernel/$l_mtype" ]; then
                module_loadable_fix
                module_loaded_fix
            fi
        else
            echo -e "\\n - module: \"$l_mname\" doesn't exist in \"$l_mdir\"\\n"
        fi
    done
    echo -e "\\n - remediation of module: \"$l_mname\" complete\\n"
}

```

References:

1. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1005, T1005.000	TA0005	M1050

1.1.2 Configure /tmp

The /tmp directory is a world-writable directory used for temporary storage by all users and some applications.

1.1.2.1 Ensure /tmp is a separate partition (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The /tmp directory is a world-writable directory used for temporary storage by all users and some applications.

Rationale:

Making /tmp its own file system allows an administrator to set additional mount options such as the `noexec` option on the mount, making /tmp useless for an attacker to install executable code. It would also prevent an attacker from establishing a hard link to a system `setuid` program and wait for it to be updated. Once the program was updated, the hard link would be broken and the attacker would have his own copy of the program. If the program happened to have a security vulnerability, the attacker could continue to exploit the known flaw.

This can be accomplished by either mounting `tmpfs` to /tmp, or creating a separate partition for /tmp.

Impact:

By design files saved to /tmp should have no expectation of surviving a reboot of the system. `tmpfs` is ram based and all files stored to `tmpfs` will be lost when the system is rebooted.

If files need to be persistent through a reboot, they should be saved to `/var/tmp` not /tmp.

Since the /tmp directory is intended to be world-writable, there is a risk of resource exhaustion if it is not bound to `tmpfs` or a separate partition. Running out of /tmp space is a problem regardless of what kind of filesystem lies under it, but in a configuration where /tmp is not a separate file system it will essentially have the whole disk available, as the default installation only creates a single / partition. On the other hand, a RAM-based /tmp (as with `tmpfs`) will almost certainly be much smaller, which can lead to applications filling up the filesystem much more easily. Another alternative is to create a dedicated partition for /tmp from a separate volume or disk. One of the downsides of a disk-based dedicated partition is that it will be slower than `tmpfs` which is RAM-based.

Audit:

Run the following command and verify the output shows that `/tmp` is mounted. Particular requirements pertaining to mount options are covered in ensuing sections.

```
# findmnt -nk /tmp
```

Example output:

```
/tmp tmpfs tmpfs rw,nosuid,nodev,noexec
```

Ensure that systemd will mount the `/tmp` partition at boot time.

```
# systemctl is-enabled tmp.mount
```

Example output:

```
generated
```

Verify output is not masked or disabled.

Note: By default systemd will output `generated` if there is an entry in `/etc/fstab` for `/tmp`. This just means systemd will use the entry in `/etc/fstab` instead of its default unit file configuration for `/tmp`.

Remediation:

First ensure that systemd is correctly configured to ensure that `/tmp` will be mounted at boot time.

```
# systemctl unmask tmp.mount
```

For specific configuration requirements of the `/tmp` mount for your environment, modify `/etc/fstab`.

Example of using `tmpfs` with specific mount options:

```
tmpfs /tmp tmpfs defaults,rw,nosuid,nodev,noexec,relatime,size=2G 0  
0
```

Example of using a volume or disk with specific mount options. The source location of the volume or disk will vary depending on your environment.

```
<device> /tmp <fstype> defaults,nodev,nosuid,noexec 0 0
```

References:

1. <https://www.freedesktop.org/wiki/Software/systemd/APIFileSystems/>
2. <https://www.freedesktop.org/software/systemd/man/systemd-fstab-generator.html>
3. NIST SP 800-53 Rev. 5: CM-7

Additional Information:

If an entry for `/tmp` exists in `/etc/fstab` it will take precedence over entries in systemd default unit file located at `/usr/lib/systemd/system/tmp.mount`.

`/tmp` utilizing `tmpfs` can be resized using the `size={size}` parameter in the relevant entry in `/etc/fstab`.

Note: In an environment where the main system is diskless and connected to iSCSI, entries in `/etc/fstab` may not take precedence.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1499, T1499.001	TA0005	M1022

1.1.2.2 Ensure nodev option set on /tmp partition (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `nodev` mount option specifies that the filesystem cannot contain special devices.

Rationale:

Since the `/tmp` filesystem is not intended to support devices, set this option to ensure that users cannot create a block or character special devices in `/tmp`.

Audit:

Verify that the `nodev` option is set for the `/tmp` mount.

Run the following command to verify that the `nodev` mount option is set.

Example:

```
# findmnt -kn /tmp | grep nodev
/tmp    tmpfs    tmpfs    rw,nosuid,nodev,noexec,relatime,seclabel
```

Remediation:

Edit the `/etc/fstab` file and add `nodev` to the fourth field (mounting options) for the `/tmp` partition.

Example:

```
<device> /tmp    <fstype>    defaults,rw,nosuid,nodev,noexec,relatime 0 0
```

Run the following command to remount `/tmp` with the configured options:

```
# mount -o remount /tmp
```

References:

1. See the `fstab(5)` manual page for more information.

Additional Information:

NIST SP 800-53 Rev. 5:

- CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1200, T1200.000	TA0005	M1022

1.1.2.3 Ensure noexec option set on /tmp partition (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `noexec` mount option specifies that the filesystem cannot contain executable binaries.

Rationale:

Since the `/tmp` filesystem is only intended for temporary file storage, set this option to ensure that users cannot run executable binaries from `/tmp`.

Audit:

Verify that the `noexec` option is set for the `/tmp` mount.

Run the following command to verify that the `noexec` mount option is set.

Example:

```
# findmnt -kn /tmp | grep noexec  
  
/tmp tmpfs tmpfs rw,nosuid,nodev,noexec,relatime,seclabel
```

Remediation:

Edit the `/etc/fstab` file and add `noexec` to the fourth field (mounting options) for the `/tmp` partition.

Example:

```
<device> /tmp <fstype> defaults,rw,nosuid,nodev,noexec,relatime 0 0
```

Run the following command to remount `/tmp` with the configured options:

```
# mount -o remount /tmp
```

References:

1. See the `fstab(5)` manual page for more information.
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1204, T1204.002	TA0005	M1022

1.1.2.4 Ensure nosuid option set on /tmp partition (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `nosuid` mount option specifies that the filesystem cannot contain `setuid` files.

Rationale:

Since the `/tmp` filesystem is only intended for temporary file storage, set this option to ensure that users cannot create `setuid` files in `/tmp`.

Audit:

Verify that the `nosuid` option is set for the `/tmp` mount.

Run the following command to verify that the `nosuid` mount option is set.

Example:

```
# findmnt -kn /tmp | grep nosuid  
  
/tmp    tmpfs    tmpfs    rw,nosuid,nodev,noexec,relatime,seclabel
```

Remediation:

Edit the `/etc/fstab` file and add `nosuid` to the fourth field (mounting options) for the `/tmp` partition.

Example:

```
<device> /tmp    <fstype>    defaults,rw,nosuid,nodev,noexec,relatime 0 0
```

Run the following command to remount `/tmp` with the configured options:

```
# mount -o remount /tmp
```

References:

1. See the `fstab(5)` manual page for more information.
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1548, T1548.001	TA0005	M1022

1.1.3 Configure /var

The `/var` directory is used by daemons and other system services to temporarily store dynamic data. Some directories created by these processes may be world-writable.

1.1.3.1 Ensure separate partition exists for /var (Automated)

Profile Applicability:

- Level 2 - Server

Description:

The `/var` directory is used by daemons and other system services to temporarily store dynamic data. Some directories created by these processes may be world-writable.

Rationale:

The reasoning for mounting `/var` on a separate partition is as follows.

Protection from resource exhaustion

The default installation only creates a single `/` partition. Since the `/var` directory may contain world-writable files and directories, there is a risk of resource exhaustion. It will essentially have the whole disk available to fill up and impact the system as a whole. In addition, other operations on the system could fill up the disk unrelated to `/var` and cause unintended behavior across the system as the disk is full. See `man auditd.conf` for details.

Fine grained control over the mount

Configuring `/var` as its own file system allows an administrator to set additional mount options such as `noexec/nosuid/nODEV`. These options limits an attackers ability to create exploits on the system. Other options allow for specific behavior. See `man mount` for exact details regarding filesystem-independent and filesystem-specific options.

Protection from exploitation

An example of exploiting `/var` may be an attacker establishing a hard-link to a system `setuid` program and wait for it to be updated. Once the program was updated, the hard-link would be broken and the attacker would have his own copy of the program. If the program happened to have a security vulnerability, the attacker could continue to exploit the known flaw.

Impact:

Resizing filesystems is a common activity in cloud-hosted servers. Separate filesystem partitions may prevent successful resizing, or may require the installation of additional tools solely for the purpose of resizing operations. The use of these additional tools may introduce their own security considerations.

Audit:

Run the following command and verify output shows `/var` is mounted.

Example:

```
# findmnt --kernel /var

TARGET SOURCE      FSTYPE OPTIONS
/var    /dev/sdb  ext4    rw,relatime,seclabel,data=ordered
```

Remediation:

For new installations, during installation create a custom partition setup and specify a separate partition for `/var`.

For systems that were previously installed, create a new partition and configure `/etc/fstab` as appropriate.

References:

1. AJ Lewis, "LVM HOWTO", <http://tldp.org/HOWTO/LVM-HOWTO/>
2. NIST SP 800-53 Rev. 5: MP-2, AC-3

Additional Information:

When modifying `/var` it is advisable to bring the system to emergency mode (so `auditd` is not running), rename the existing directory, mount the new file system, and migrate the data over before returning to multi-user mode.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1499, T1499.001	TA0006	M1022

1.1.3.2 Ensure nodev option set on /var partition (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `nodev` mount option specifies that the filesystem cannot contain special devices.

Rationale:

Since the `/var` filesystem is not intended to support devices, set this option to ensure that users cannot create a block or character special devices in `/var`.

Audit:

Verify that the `nodev` option is set for the `/var` mount.

Run the following command to verify that the `nodev` mount option is set.

Example:

```
# findmnt --kernel /var | grep nodev  
  
/var    /dev/sdb ext4    rw,nosuid,nodev,noexec,relatime,seclabel
```

Remediation:

Edit the `/etc/fstab` file and add `nodev` to the fourth field (mounting options) for the `/var` partition.

Example:

```
<device> /var    <fstype>        defaults,rw,nosuid,nodev,noexec,relatime 0 0
```

Run the following command to remount `/var` with the configured options:

```
# mount -o remount /var
```

References:

1. See the `fstab(5)` manual page for more information.
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1200, T1200.000	TA0005	M1022

1.1.3.3 Ensure nosuid option set on /var partition (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `nosuid` mount option specifies that the filesystem cannot contain `setuid` files.

Rationale:

Since the `/var` filesystem is only intended for variable files such as logs, set this option to ensure that users cannot create `setuid` files in `/var`.

Audit:

Verify that the `nosuid` option is set for the `/var` mount.

Run the following command to verify that the `nosuid` mount option is set.

Example:

```
# findmnt --kernel /var | grep nosuid  
  
/var    /dev/sdb ext4    rw,nosuid,nodev,noexec,relatime,seclabel
```

Remediation:

Edit the `/etc/fstab` file and add `nosuid` to the fourth field (mounting options) for the `/var` partition.

Example:

```
<device> /var    <fstype>    defaults,rw,nosuid,nodev,noexec,relatime 0 0
```

Run the following command to remount `/var` with the configured options:

```
# mount -o remount /var
```

References:

1. See the `fstab(5)` manual page for more information.
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1548, T1548.001	TA0005	M1022

1.1.4 Configure `/var/tmp`

The `/var/tmp` directory is a world-writable directory used for temporary storage by all users and some applications. Temporary files residing in `/var/tmp` are to be preserved between reboots.

1.1.4.1 Ensure separate partition exists for /var/tmp (Automated)

Profile Applicability:

- Level 2 - Server

Description:

The `/var/tmp` directory is a world-writable directory used for temporary storage by all users and some applications. Temporary files residing in `/var/tmp` are to be preserved between reboots.

Rationale:

The reasoning for mounting `/var/tmp` on a separate partition is as follows.

Protection from resource exhaustion

The default installation only creates a single `/` partition. Since the `/var/tmp` directory may contain world-writable files and directories, there is a risk of resource exhaustion. It will essentially have the whole disk available to fill up and impact the system as a whole. In addition, other operations on the system could fill up the disk unrelated to `/var/tmp` and cause potential disruption to daemons as the disk is full.

Fine grained control over the mount

Configuring `/var/tmp` as its own file system allows an administrator to set additional mount options such as `noexec/nosuid/nODEV`. These options limits an attackers ability to create exploits on the system. Other options allow for specific behavior. See `man mount` for exact details regarding filesystem-independent and filesystem-specific options.

Protection from exploitation

An example of exploiting `/var/tmp` may be an attacker establishing a hard-link to a system `setuid` program and wait for it to be updated. Once the program was updated, the hard-link would be broken and the attacker would have his own copy of the program. If the program happened to have a security vulnerability, the attacker could continue to exploit the known flaw.

Impact:

Resizing filesystems is a common activity in cloud-hosted servers. Separate filesystem partitions may prevent successful resizing, or may require the installation of additional tools solely for the purpose of resizing operations. The use of these additional tools may introduce their own security considerations.

Audit:

Run the following command and verify output shows `/var/tmp` is mounted.

Example:

```
# findmnt --kernel /var/tmp

TARGET SOURCE      FSTYPE OPTIONS
/var/tmp /dev/sdb ext4    rw,relatime,seclabel,data=ordered
```

Remediation:

For new installations, during installation create a custom partition setup and specify a separate partition for `/var/tmp`.

For systems that were previously installed, create a new partition and configure `/etc/fstab` as appropriate.

References:

1. AJ Lewis, "LVM HOWTO", <http://tldp.org/HOWTO/LVM-HOWTO/>
2. NIST SP 800-53 Rev. 5: CM-6

Additional Information:

When modifying `/var/tmp` it is advisable to bring the system to emergency mode (so `auditd` is not running), rename the existing directory, mount the new file system, and migrate the data over before returning to multi-user mode.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1499, T1499.001	TA0005	M1022

1.1.4.2 Ensure noexec option set on /var/tmp partition (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `noexec` mount option specifies that the filesystem cannot contain executable binaries.

Rationale:

Since the `/var/tmp` filesystem is only intended for temporary file storage, set this option to ensure that users cannot run executable binaries from `/var/tmp`.

Audit:

Verify that the `noexec` option is set for the `/var/tmp` mount.

Run the following command to verify that the `noexec` mount option is set.

Example:

```
# findmnt --kernel /var/tmp | grep noexec  
  
/var/tmp    /dev/sdb ext4    rw,nosuid,nodev,noexec,relatime,seclabel
```

Remediation:

Edit the `/etc/fstab` file and add `noexec` to the fourth field (mounting options) for the `/var/tmp` partition.

Example:

```
<device> /var/tmp    <fstype>    defaults,rw,nosuid,nodev,noexec,relatime 0  
0
```

Run the following command to remount `/var/tmp` with the configured options:

```
# mount -o remount /var/tmp
```

References:

1. See the `fstab(5)` manual page for more information.
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1204, T1204.002	TA0005	M1022

1.1.4.3 Ensure nosuid option set on /var/tmp partition (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `nosuid` mount option specifies that the filesystem cannot contain `setuid` files.

Rationale:

Since the `/var/tmp` filesystem is only intended for temporary file storage, set this option to ensure that users cannot create `setuid` files in `/var/tmp`.

Audit:

Verify that the `nosuid` option is set for the `/var/tmp` mount.

Run the following command to verify that the `nosuid` mount option is set.

Example:

```
# findmnt --kernel /var/tmp | grep nosuid  
  
/var/tmp    /dev/sdb ext4    rw,nosuid,nodev,noexec,relatime,seclabel
```

Remediation:

Edit the `/etc/fstab` file and add `nosuid` to the fourth field (mounting options) for the `/var/tmp` partition.

Example:

```
<device> /var/tmp    <fstype>    defaults,rw,nosuid,nodev,noexec,relatime 0  
0
```

Run the following command to remount `/var/tmp` with the configured options:

```
# mount -o remount /var/tmp
```

References:

1. See the `fstab(5)` manual page for more information.
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1548, T1548.001	TA0005	M1022

1.1.4.4 Ensure nodev option set on /var/tmp partition (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `nodev` mount option specifies that the filesystem cannot contain special devices.

Rationale:

Since the `/var/tmp` filesystem is not intended to support devices, set this option to ensure that users cannot create a block or character special devices in `/var/tmp`.

Audit:

Verify that the `nodev` option is set for the `/var/tmp` mount.

Run the following command to verify that the `nodev` mount option is set.

Example:

```
# findmnt --kernel /var/tmp | grep nodev
/var/tmp    /dev/sdb ext4    rw,nosuid,nodev,noexec,relatime,seclabel
```

Remediation:

Edit the `/etc/fstab` file and add `nodev` to the fourth field (mounting options) for the `/var/tmp` partition.

Example:

```
<device> /var/tmp    <fstype>    defaults,rw,nosuid,nodev,noexec,relatime 0
0
```

Run the following command to remount `/var/tmp` with the configured options:

```
# mount -o remount /var/tmp
```

References:

1. See the `fstab(5)` manual page for more information.
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1499, T1499.001	TA0005	M1022

1.1.5 Configure `/var/log`

The `/var/log` directory is used by system services to store log data.

1.1.5.1 Ensure separate partition exists for /var/log (Automated)

Profile Applicability:

- Level 2 - Server

Description:

The `/var/log` directory is used by system services to store log data.

Rationale:

The reasoning for mounting `/var/log` on a separate partition is as follows.

Protection from resource exhaustion

The default installation only creates a single `/` partition. Since the `/var/log` directory contains log files which can grow quite large, there is a risk of resource exhaustion. It will essentially have the whole disk available to fill up and impact the system as a whole.

Fine grained control over the mount

Configuring `/var/log` as its own file system allows an administrator to set additional mount options such as `noexec/nosuid/nODEV`. These options limit an attackers ability to create exploits on the system. Other options allow for specific behavior. See `man mount` for exact details regarding filesystem-independent and filesystem-specific options.

Protection of log data

As `/var/log` contains log files, care should be taken to ensure the security and integrity of the data and mount point.

Impact:

Resizing filesystems is a common activity in cloud-hosted servers. Separate filesystem partitions may prevent successful resizing, or may require the installation of additional tools solely for the purpose of resizing operations. The use of these additional tools may introduce their own security considerations.

Audit:

Run the following command and verify output shows `/var/log` is mounted:

```
# findmnt --kernel /var/log

TARGET          SOURCE      FSTYPE OPTIONS
/var/log /dev/sdb ext4    rw,relatime,seclabel,data=ordered
```

Remediation:

For new installations, during installation create a custom partition setup and specify a separate partition for `/var/log`.

For systems that were previously installed, create a new partition and configure `/etc/fstab` as appropriate.

References:

1. AJ Lewis, "LVM HOWTO", <http://tldp.org/HOWTO/LVM-HOWTO/>
2. NIST SP 800-53 Rev. 5: CM-6

Additional Information:

When modifying `/var/log` it is advisable to bring the system to emergency mode (so `auditd` is not running), rename the existing directory, mount the new file system, and migrate the data over before returning to multiuser mode.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.3 <u>Ensure Adequate Audit Log Storage</u> Ensure that logging destinations maintain adequate storage to comply with the enterprise's audit log management process.			
v7	6.4 <u>Ensure adequate storage for logs</u> Ensure that all systems that store logs have adequate storage space for the logs generated.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1499, T1499.001	TA0005	M1022

1.1.5.2 Ensure nodev option set on /var/log partition (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `nodev` mount option specifies that the filesystem cannot contain special devices.

Rationale:

Since the `/var/log` filesystem is not intended to support devices, set this option to ensure that users cannot create a block or character special devices in `/var/log`.

Audit:

Verify that the `nodev` option is set for the `/var/log` mount.

Run the following command to verify that the `nodev` mount option is set.

Example:

```
# findmnt --kernel /var/log | grep nodev
/var/log    /dev/sdb ext4    rw,nosuid,nodev,noexec,relatime,seclabel
```

Remediation:

Edit the `/etc/fstab` file and add `nodev` to the fourth field (mounting options) for the `/var/log` partition.

Example:

```
<device> /var/log    <fstype>    defaults,rw,nosuid,nodev,noexec,relatime 0
0
```

Run the following command to remount `/var/log` with the configured options:

```
# mount -o remount /var/log
```

References:

1. See the `fstab(5)` manual page for more information.
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1200, T1200.000	TA0005	M1038

1.1.5.3 Ensure noexec option set on /var/log partition (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `noexec` mount option specifies that the filesystem cannot contain executable binaries.

Rationale:

Since the `/var/log` filesystem is only intended for log files, set this option to ensure that users cannot run executable binaries from `/var/log`.

Audit:

Verify that the `noexec` option is set for the `/var/log` mount.

Run the following command to verify that the `noexec` mount option is set.

Example:

```
# findmnt --kernel /var/log | grep noexec
/var/log    /dev/sdb ext4  rw,nosuid,nodev,noexec,relatime,seclabel
```

Remediation:

Edit the `/etc/fstab` file and add `noexec` to the fourth field (mounting options) for the `/var/log` partition.

Example:

```
<device> /var/log    <fstype>    defaults,rw,nosuid,nodev,noexec,relatime 0
0
```

Run the following command to remount `/var/log` with the configured options:

```
# mount -o remount /var/log
```

References:

1. See the `fstab(5)` manual page for more information.
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1204, T1204.002	TA0005	M1022

1.1.5.4 Ensure nosuid option set on /var/log partition (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `nosuid` mount option specifies that the filesystem cannot contain `setuid` files.

Rationale:

Since the `/var/log` filesystem is only intended for log files, set this option to ensure that users cannot create `setuid` files in `/var/log`.

Audit:

Verify that the `nosuid` option is set for the `/var/log` mount.

Run the following command to verify that the `nosuid` mount option is set.

Example:

```
# findmnt --kernel /var/log | grep nosuid  
  
/var/log    /dev/sdb ext4    rw,nosuid,nodev,noexec,relatime,seclabel
```

Remediation:

Edit the `/etc/fstab` file and add `nosuid` to the fourth field (mounting options) for the `/var/log` partition.

Example:

```
<device> /var/log    <fstype>    defaults,rw,nosuid,nodev,noexec,relatime 0  
0
```

Run the following command to remount `/var/log` with the configured options:

```
# mount -o remount /var/log
```

References:

1. See the `fstab(5)` manual page for more information.
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1548, T1548.001	TA0005	M1022

1.1.6 Configure `/var/log/audit`

The auditing daemon, `auditd`, stores log data in the `/var/log/audit` directory.

1.1.6.1 Ensure separate partition exists for /var/log/audit (Automated)

Profile Applicability:

- Level 2 - Server

Description:

The auditing daemon, `auditd`, stores log data in the `/var/log/audit` directory.

Rationale:

The reasoning for mounting `/var/log/audit` on a separate partition is as follows.

Protection from resource exhaustion

The default installation only creates a single `/` partition. Since the `/var/log/audit` directory contains the `audit.log` file which can grow quite large, there is a risk of resource exhaustion. It will essentially have the whole disk available to fill up and impact the system as a whole. In addition, other operations on the system could fill up the disk unrelated to `/var/log/audit` and cause `auditd` to trigger its `space_left_action` as the disk is full. See `man auditd.conf` for details.

Fine grained control over the mount

Configuring `/var/log/audit` as its own file system allows an administrator to set additional mount options such as `noexec/nosuid/nodev`. These options limit an attacker's ability to create exploits on the system. Other options allow for specific behavior. See `man mount` for exact details regarding filesystem-independent and filesystem-specific options.

Protection of audit data

As `/var/log/audit` contains audit logs, care should be taken to ensure the security and integrity of the data and mount point.

Impact:

Resizing filesystems is a common activity in cloud-hosted servers. Separate filesystem partitions may prevent successful resizing, or may require the installation of additional tools solely for the purpose of resizing operations. The use of these additional tools may introduce their own security considerations.

Audit:

Run the following command and verify output shows `/var/log/audit` is mounted:

```
# findmnt --kernel /var/log/audit

TARGET          SOURCE      FSTYPE OPTIONS
/var/log/audit  /dev/sdb   ext4    rw,relatime,seclabel,data=ordered
```

Remediation:

For new installations, during installation create a custom partition setup and specify a separate partition for `/var/log/audit`.

For systems that were previously installed, create a new partition and configure `/etc/fstab` as appropriate.

References:

1. AJ Lewis, "LVM HOWTO", <http://tldp.org/HOWTO/LVM-HOWTO/>
2. NIST SP 800-53 Rev. 5: CM-6

Additional Information:

When modifying `/var/log/audit` it is advisable to bring the system to emergency mode (so `auditd` is not running), rename the existing directory, mount the new file system, and migrate the data over before returning to multi-user mode.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.3 <u>Ensure Adequate Audit Log Storage</u> Ensure that logging destinations maintain adequate storage to comply with the enterprise's audit log management process.			
v7	6.4 <u>Ensure adequate storage for logs</u> Ensure that all systems that store logs have adequate storage space for the logs generated.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1499, T1499.001	TA0005	M1022

1.1.6.2 Ensure noexec option set on /var/log/audit partition (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `noexec` mount option specifies that the filesystem cannot contain executable binaries.

Rationale:

Since the `/var/log/audit` filesystem is only intended for audit logs, set this option to ensure that users cannot run executable binaries from `/var/log/audit`.

Audit:

Verify that the `noexec` option is set for the `/var/log/audit` mount.
Run the following command to verify that the `noexec` mount option is set.
Example:

```
# findmnt --kernel /var/log/audit | grep noexec  
  
/var/log/audit /dev/sdb ext4 rw,nosuid,nodev,noexec,relatime,seclabel
```

Remediation:

Edit the `/etc/fstab` file and add `noexec` to the fourth field (mounting options) for the `/var` partition.
Example:

```
<device> /var/log/audit <fstype>  
defaults,rw,nosuid,nodev,noexec,relatime 0 0
```

Run the following command to remount `/var/log/audit` with the configured options:

```
# mount -o remount /var/log/audit
```

References:

1. See the `fstab(5)` manual page for more information.
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1204, T1204.002	TA0005	M1022

1.1.6.3 Ensure nodev option set on /var/log/audit partition (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `nodev` mount option specifies that the filesystem cannot contain special devices.

Rationale:

Since the `/var/log/audit` filesystem is not intended to support devices, set this option to ensure that users cannot create a block or character special devices in `/var/log/audit`.

Audit:

Verify that the `nodev` option is set for the `/var/log/audit` mount.
Run the following command to verify that the `nodev` mount option is set.
Example:

```
# findmnt --kernel /var/log/audit | grep nodev  
  
/var/log/audit /dev/sdb ext4 rw,nosuid,nodev,noexec,relatime,seclabel
```

Remediation:

Edit the `/etc/fstab` file and add `nodev` to the fourth field (mounting options) for the `/var/log/audit` partition.

Example:

```
<device> /var/log/audit <fstype>  
defaults,rw,nosuid,nodev,noexec,relatime 0 0
```

Run the following command to remount `/var/log/audit` with the configured options:

```
# mount -o remount /var/log/audit
```

References:

1. See the `fstab(5)` manual page for more information.
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1200, T1200.000	TA0005	M1022

1.1.6.4 Ensure nosuid option set on /var/log/audit partition (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `nosuid` mount option specifies that the filesystem cannot contain `setuid` files.

Rationale:

Since the `/var/log/audit` filesystem is only intended for variable files such as logs, set this option to ensure that users cannot create `setuid` files in `/var/log/audit`.

Audit:

Verify that the `nosuid` option is set for the `/var/log/audit` mount.
Run the following command to verify that the `nosuid` mount option is set.
Example:

```
# findmnt --kernel /var/log/audit | grep nosuid  
  
/var/log/audit /dev/sdb ext4 rw,nosuid,nodev,noexec,relatime,seclabel
```

Remediation:

Edit the `/etc/fstab` file and add `nosuid` to the fourth field (mounting options) for the `/var/log/audit` partition.

Example:

```
<device> /var/log/audit <fstype>  
defaults,rw,nosuid,nodev,noexec,relatime 0 0
```

Run the following command to remount `/var/log/audit` with the configured options:

```
# mount -o remount /var/log/audit
```

References:

1. See the `fstab(5)` manual page for more information.
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1548, T1548.001	TA0005	M1022

1.1.7 Configure /home

Please note that home directories could be mounted anywhere and are not necessarily restricted to `/home`, nor restricted to a single location, nor is the name restricted in any way.

Checks can be made by looking in `/etc/passwd`, looking over the mounted file systems with `mount` or querying the relevant database with `getent`.

1.1.7.1 Ensure separate partition exists for /home (Automated)

Profile Applicability:

- Level 2 - Server

Description:

The `/home` directory is used to support disk storage needs of local users.

Rationale:

The reasoning for mounting `/home` on a separate partition is as follows.

Protection from resource exhaustion

The default installation only creates a single `/` partition. Since the `/home` directory contains user generated data, there is a risk of resource exhaustion. It will essentially have the whole disk available to fill up and impact the system as a whole. In addition, other operations on the system could fill up the disk unrelated to `/home` and impact all local users.

Fine grained control over the mount

Configuring `/home` as its own file system allows an administrator to set additional mount options such as `noexec/nosuid/nODEV`. These options limit an attacker's ability to create exploits on the system. In the case of `/home` options such as `usrquota/grpquota` may be considered to limit the impact that users can have on each other with regards to disk resource exhaustion. Other options allow for specific behavior. See `man mount` for exact details regarding filesystem-independent and filesystem-specific options.

Protection of user data

As `/home` contains user data, care should be taken to ensure the security and integrity of the data and mount point.

Impact:

Resizing filesystems is a common activity in cloud-hosted servers. Separate filesystem partitions may prevent successful resizing, or may require the installation of additional tools solely for the purpose of resizing operations. The use of these additional tools may introduce their own security considerations.

Audit:

Run the following command and verify output shows `/home` is mounted:

```
# findmnt --kernel /home

TARGET SOURCE     FSTYPE OPTIONS
/home   /dev/sdb ext4    rw,relatime,seclabel
```

Remediation:

For new installations, during installation create a custom partition setup and specify a separate partition for `/home`.

For systems that were previously installed, create a new partition and configure `/etc/fstab` as appropriate.

References:

1. AJ Lewis, "LVM HOWTO", <http://tldp.org/HOWTO/LVM-HOWTO/>
2. NIST SP 800-53 Rev. 5: CM-6

Additional Information:

When modifying `/home` it is advisable to bring the system to emergency mode (so `auditd` is not running), rename the existing directory, mount the new file system, and migrate the data over before returning to multi-user mode.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1499, T1499.001	TA0005	M1022

1.1.7.2 Ensure nodev option set on /home partition (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `nodev` mount option specifies that the filesystem cannot contain special devices.

Rationale:

Since the `/home` filesystem is not intended to support devices, set this option to ensure that users cannot create a block or character special devices in `/var`.

Audit:

Verify that the `nodev` option is set for the `/home` mount.

Run the following command to verify that the `nodev` mount option is set.

Example:

```
# findmnt --kernel /home | grep nodev
/home      /dev/sdb ext4  rw,nosuid,nodev,noexec,relatime,seclabel
```

Remediation:

Edit the `/etc/fstab` file and add `nodev` to the fourth field (mounting options) for the `/home` partition.

Example:

```
<device> /home    <fstype>          defaults,rw,nosuid,nodev,noexec,relatime 0 0
```

Run the following command to remount `/home` with the configured options:

```
# mount -o remount /home
```

References:

1. See the `fstab(5)` manual page for more information.
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1200, T1200.000	TA0005	M1038

1.1.7.3 Ensure nosuid option set on /home partition (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `nosuid` mount option specifies that the filesystem cannot contain `setuid` files.

Rationale:

Since the `/home` filesystem is only intended for user file storage, set this option to ensure that users cannot create `setuid` files in `/home`.

Audit:

Verify that the `nosuid` option is set for the `/home` mount.

Run the following command to verify that the `nosuid` mount option is set.

Example:

```
# findmnt --kernel /home | grep nosuid  
  
/home    /dev/sdb ext4    rw,nosuid,nodev,noexec,relatime,seclabel
```

Remediation:

Edit the `/etc/fstab` file and add `nosuid` to the fourth field (mounting options) for the `/home` partition.

Example:

```
<device> /home    <fstype>      defaults,rw,nosuid,nodev,noexec,relatime 0 0
```

Run the following command to remount `/home` with the configured options:

```
# mount -o remount /home
```

References:

1. See the `fstab(5)` manual page for more information.
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1548, T1548.001	TA0005	M1022

1.1.8 Configure /dev/shm

The `/dev/shm` directory is a world-writable directory that can function as shared memory that facilitates inter process communication (IPC)

1.1.8.1 Ensure /dev/shm is a separate partition (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `/dev/shm` directory is a world-writable directory that can function as shared memory that facilitates inter process communication (IPC).

Rationale:

Making `/dev/shm` its own file system allows an administrator to set additional mount options such as the `noexec` option on the mount, making `/dev/shm` useless for an attacker to install executable code. It would also prevent an attacker from establishing a hard link to a system `setuid` program and wait for it to be updated. Once the program was updated, the hard link would be broken and the attacker would have his own copy of the program. If the program happened to have a security vulnerability, the attacker could continue to exploit the known flaw.

This can be accomplished by mounting `tmpfs` to `/dev/shm`.

Impact:

Since the `/dev/shm` directory is intended to be world-writable, there is a risk of resource exhaustion if it is not bound to a separate partition.

`/dev/shm` utilizing `tmpfs` can be resized using the `size={size}` parameter in the relevant entry in `/etc/fstab`.

Audit:

IF `/dev/shm` is to be used on the system, run the following command and verify the output shows that `/dev/shm` is mounted. Particular requirements pertaining to mount options are covered in ensuing sections.

```
# findmnt --kernel /dev/shm

TARGET SOURCE FSTYPE OPTIONS
/dev/shm tmpfs tmpfs rw,nosuid,nodev,noexec,relatime,seclabel
```

Remediation:

For specific configuration requirements of the `/dev/shm` mount for your environment, modify `/etc/fstab`.

Example of using `tmpfs` with specific mount options:

```
tmpfs    /dev/shm          tmpfs
defaults,rw,nosuid,nodev,noexec,relatime,size=2G 0 0
```

References:

1. <https://www.freedesktop.org/wiki/Software/systemd/APIFileSystems/>
2. <https://www.freedesktop.org/software/systemd/man/systemd-fstab-generator.html>
3. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 <u>Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	9.2 <u>Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1499, T1499.001	TA0005	M1022

1.1.8.2 Ensure nodev option set on /dev/shm partition (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `nodev` mount option specifies that the filesystem cannot contain special devices.

Rationale:

Since the `/dev/shm` filesystem is not intended to support devices, set this option to ensure that users cannot attempt to create special devices in `/dev/shm` partitions.

Audit:

Verify that the `nodev` option is set if a `/dev/shm` partition exists.
Run the following command and verify that nothing is returned:

```
# mount | grep -E '\s/dev/shm\s' | grep -v nodev
```

Remediation:

Edit the `/etc/fstab` file and add `nodev` to the fourth field (mounting options) for the `/dev/shm` partition. See the `fstab(5)` manual page for more information.
Run the following command to remount `/dev/shm` using the updated options from `/etc/fstab`:

```
# mount -o remount /dev/shm
```

References:

1. NIST SP 800-53 Rev. 5: AC-3, MP-2

Additional Information:

Some distributions mount `/dev/shm` through other means and require `/dev/shm` to be added to `/etc/fstab` even though it is already being mounted on boot. Others may configure `/dev/shm` in other locations and may override `/etc/fstab` configuration. Consult the documentation appropriate for your distribution.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1200, T1200.000	TA0005	M1038

1.1.8.3 Ensure noexec option set on /dev/shm partition (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `noexec` mount option specifies that the filesystem cannot contain executable binaries.

Rationale:

Setting this option on a file system prevents users from executing programs from shared memory. This deters users from introducing potentially malicious software on the system.

Audit:

Verify that the `noexec` option is set for the `/dev/shm` mount.

Run the following command to verify that the `noexec` mount option is set.

Example:

```
# findmnt --kernel /dev/shm | grep noexec  
  
/dev/shm tmpfs tmpfs rw,nosuid,nodev,noexec,relatime,seclabel
```

Remediation:

Edit the `/etc/fstab` file and add `noexec` to the fourth field (mounting options) for the `/dev/shm` partition.

Example:

```
<device> /dev/shm <fstype> defaults,rw,nosuid,nodev,noexec,relatime 0  
0
```

Run the following command to remount `/dev/shm` with the configured options:

```
# mount -o remount /dev/shm
```

NOTE It is recommended to use `tmpfs` as the device/filesystem type as `/dev/shm` is used as shared memory space by applications.

References:

1. See the `fstab(5)` manual page for more information.
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1204, T1204.002	TA0005	M1022

1.1.8.4 Ensure nosuid option set on /dev/shm partition (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `nosuid` mount option specifies that the filesystem cannot contain `setuid` files.

Rationale:

Setting this option on a file system prevents users from introducing privileged programs onto the system and allowing non-root users to execute them.

Audit:

Verify that the `nosuid` option is set if a `/dev/shm` partition exists.
Run the following command and verify that nothing is returned:

```
# mount | grep -E '\s/dev/shm\s' | grep -v nosuid
```

Remediation:

Edit the `/etc/fstab` file and add `nosuid` to the fourth field (mounting options) for the `/dev/shm` partition. See the `fstab(5)` manual page for more information.
Run the following command to remount `/dev/shm` using the updated options from `/etc/fstab`:

```
# mount -o remount /dev/shm
```

References:

1. NIST SP 800-53 Rev. 5: AC-3, MP-2

Additional Information:

Some distributions mount `/dev/shm` through other means and require `/dev/shm` to be added to `/etc/fstab` even though it is already being mounted on boot. Others may configure `/dev/shm` in other locations and may override `/etc/fstab` configuration. Consult the documentation appropriate for your distribution.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1548, T1548.001	TA0005	M1022

1.1.9 Ensure usb-storage is disabled (Automated)

Profile Applicability:

- Level 1 - Server

Description:

USB storage provides a means to transfer and store files ensuring persistence and availability of the files independent of network connection status. Its popularity and utility has led to USB-based malware being a simple and common means for network infiltration and a first step to establishing a persistent threat within a networked environment.

Rationale:

Restricting USB access on the system will decrease the physical attack surface for a device and diminish the possible vectors to introduce malware.

Impact:

Disabling the `usb-storage` module will disable any usage of USB storage devices.

If requirements and local site policy allow the use of such devices, other solutions should be configured accordingly instead. One example of a commonly used solution is USBGuard.

Audit:

Run the following script to verify `usb-storage` is disabled:

```
#!/usr/bin/env bash

{
    l_output="" l_output2=""
    l_mname="usb-storage" # set module name
    # Check if the module exists on the system
    if [ -z "$(modprobe -n -v "$l_mname" 2>&1 | grep -Pi --
"\h*modprobe:\h+FATAL:\h+Module\h+$l_mname\h+not\h+found\h+in\h+directory)" )"
]; then
        # Check how module will be loaded
        l_loadable="$(modprobe -n -v "$l_mname")"
        [ "$(wc -l <<< "$l_loadable")" -gt "1" ] && l_loadable="$(grep -P --
"(\h*install|\b$l_mname)\b" <<< "$l_loadable")"
        if grep -Pq -- '\h*install \/\bin\/(true|false)' <<< "$l_loadable";
then
            l_output="$l_output\n - module: \"$l_mname\" is not loadable:
\"$l_loadable\"
        else
            l_output2="$l_output2\n - module: \"$l_mname\" is loadable:
\"$l_loadable\"
        fi
        # Check is the module currently loaded
        if ! lsmod | grep "$l_mname" > /dev/null 2>&1; then
            l_output="$l_output\n - module: \"$l_mname\" is not loaded"
        else
            l_output2="$l_output2\n - module: \"$l_mname\" is loaded"
        fi
        # Check if the module is deny listed
        if modprobe --showconfig | grep -Pq -- "\h*blacklist\h+$(tr '-' '_'
<<< "$l_mname")\b"; then
            l_output="$l_output\n - module: \"$l_mname\" is deny listed in:
\"$(grep -Pl -- "\h*blacklist\h+$l_mname\b" /etc/modprobe.d/*)\n"
        else
            l_output2="$l_output2\n - module: \"$l_mname\" is not deny listed"
        fi
    else
        l_output="$l_output\n - Module \"$l_mname\" doesn't exist on the
system"
    fi
    # Report results. If no failures output in l_output2, we pass
    if [ -z "$l_output2" ]; then
        echo -e "\n- Audit Result:\n  ** PASS **\n$l_output\n"
    else
        echo -e "\n- Audit Result:\n  ** FAIL **\n - Reason(s) for audit
failure:\n$l_output2\n"
        [ -n "$l_output" ] && echo -e "\n- Correctly set:\n$l_output\n"
    fi
}
```

Remediation:

Run the following script to disable usb-storage:

```
#!/usr/bin/env bash

{
    l_mname="usb-storage" # set module name
    # Check if the module exists on the system
    if [ -z "$(modprobe -n -v "$l_mname" 2>&1 | grep -Pi --
"\h*modprobe:\h+FATAL:\h+Module\h+$l_mname\h+not\h+found\h+in\h+directory")"
]; then
        # Remediate loadable
        l_loadable="$(modprobe -n -v "$l_mname")"
        [ "$(wc -l <<< "$l_loadable")" -gt "1" ] && l_loadable="$(grep -P --
"(\h*install|\b$l_mname)\b" <<< "$l_loadable")"
        if ! grep -Pq -- '^\h*install \/\bin\/(true|false)' <<< "$l_loadable";
then
            echo -e " - setting module: \"$l_mname\" to be not loadable"
            echo -e "install $l_mname /bin/false" >>
/etc/modprobe.d/"$l_mname".conf
        fi
        # Remediate loaded
        if lsmod | grep "$l_mname" > /dev/null 2>&1; then
            echo -e " - unloading module \"$l_mname\""
            modprobe -r "$l_mname"
        fi
        # Remediate deny list
        if ! modprobe --showconfig | grep -Pq -- "^\h*blacklist\h+$(tr '-' '_'
<<< "$l_mname")\b"; then
            echo -e " - deny listing \"$l_mname\""
            echo -e "blacklist $l_mname" >> /etc/modprobe.d/"$l_mname".conf
        fi
    else
        echo -e " - Nothing to remediate\n - Module \"$l_mname\" doesn't exist
on the system"
    fi
}
```


References:

1. NIST SP 800-53 Rev. 5: SI-3

Additional Information:

An alternative solution to disabling the usb-storage module may be found in USBGuard. Use of USBGuard and construction of USB device policies should be done in alignment with site policy.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	10.3 <u>Disable Autorun and Autoplay for Removable Media</u> Disable autorun and autoplay auto-execute functionality for removable media.			
v7	13.7 <u>Manage USB Devices</u> If USB storage devices are required, enterprise software should be used that can configure systems to allow the use of specific devices. An inventory of such devices should be maintained.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1052, T1052.001, T1091, T1091.000, T1200, T1200.000	TA0001, TA0010	M1034

1.2 Configure Software and Patch Management

Fedora 34/CentOS 9 stream derived Linux distributions use dnf (previously yum) to install and update software packages. Patch management procedures may vary widely between enterprises. Large enterprises may choose to install a local updates server that can be used in place of their distributions servers, whereas a single deployment of a system may prefer to get updates directly. Updates can be performed automatically or manually, depending on the site's policy for patch management. Organizations may prefer to test patches against their environment on a non-production system before rolling out to production.

Outdated software is vulnerable to cyber criminals and hackers. Software updates help reduce the risk to your organization. The release of software update notes often reveal the patched exploitable entry points to the public. Public knowledge of these exploits can make your organization more vulnerable to malicious actors attempting to gain entry to your system's data.

Software updates often offer new and improved features and speed enhancements

For the purpose of this benchmark, the requirement is to ensure that a patch management process is defined and maintained, the specifics of which are left to the organization.

1.2.1 Ensure GPG keys are configured (Manual)

Profile Applicability:

- Level 1 - Server

Description:

The RPM Package Manager implements GPG key signing to verify package integrity during and after installation.

Rationale:

It is important to ensure that updates are obtained from a valid source to protect against spoofing that could lead to the inadvertent installation of malware on the system. To this end, verify that GPG keys are configured correctly for your system.

Audit:

List all GPG key URLs

Each repository should have a `gpgkey` with a URL pointing to the location of the GPG key, either local or remote.

```
# grep -r gpgkey /etc/yum.repos.d/* /etc/dnf/dnf.conf
```

List installed GPG keys

Run the following command to list the currently installed keys. These are the active keys used for verification and installation of RPMs. The packages are fake, they are generated on the fly by `dnf` or `rpm` during the import of keys from the URL specified in the repository configuration.

Example:

```
# for RPM_PACKAGE in $(rpm -q gpg-pubkey); do
    echo "RPM: ${RPM_PACKAGE}"
    RPM_SUMMARY=$(rpm -q --queryformat "%{SUMMARY}" "${RPM_PACKAGE}")
    RPM_PACKAGER=$(rpm -q --queryformat "%{PACKAGER}" "${RPM_PACKAGE}")
    RPM_DATE=$(date +%Y-%m-%d -d "1970-1-1+$(0x$(rpm -q --queryformat
"%{RELEASE}" "${RPM_PACKAGE}") ))sec")
    RPM_KEY_ID=$(rpm -q --queryformat "%{VERSION}" "${RPM_PACKAGE}")
    echo "Packager: ${RPM_PACKAGER}
Summary: ${RPM_SUMMARY}
Creation date: ${RPM_DATE}
Key ID: ${RPM_KEY_ID}
"
done

RPM: gpg-pubkey-9db62fb1-59920156
Packager: Fedora 28 (28) <fedora-28@fedoraproject.org>
Summary: gpg(Fedora 28 (28) <fedora-28@fedoraproject.org>)
Creation date: 2017-08-14
Key ID: 9db62fb1

RPM: gpg-pubkey-09eab3f2-595fbba3
Packager: RPM Fusion free repository for Fedora (28) <rpmfusion-
buildsys@lists.rpmfusion.org>
Summary: gpg(RPM Fusion free repository for Fedora (28) <rpmfusion-
buildsys@lists.rpmfusion.org>)
Creation date: 2017-07-07
Key ID: 09eab3f2
```

The format of the package (gpg-pubkey-9db62fb1-59920156) is important to understand for verification. Using the above example, it consists of three parts:

1. The general prefix name for all imported GPG keys: gpg-pubkey-
2. The version, which is the GPG key ID: 9db62fb1
3. The release is the date of the key in UNIX timestamp in hexadecimal: 59920156

With both the date and the GPG key ID, check the relevant repositories public key page to confirm that the keys are indeed correct.

Query locally available GPG keys

Repositories that store their respective GPG keys on disk should do so in /etc/pki/rpm-gpg/. These keys are available for immediate import either when dnf is asked to install a relevant package from the repository or when an administrator imports the key directly with the rpm --import command.

To find where these keys come from run:

```
# for PACKAGE in $(find /etc/pki/rpm-gpg/ -type f -exec rpm -qf {} \; | sort
-u); do rpm -q --queryformat "%{NAME}-%{VERSION} %{PACKAGER} %{SUMMARY}\\n"
"${PACKAGE}"; done
```

Remediation:

Update your package manager GPG keys in accordance with site policy.

References:

1. NIST SP 800-53 Rev. 5: SI-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>7.3 Perform Automated Operating System Patch Management</u> Perform operating system updates on enterprise assets through automated patch management on a monthly, or more frequent, basis.			
v8	<u>7.4 Perform Automated Application Patch Management</u> Perform application updates on enterprise assets through automated patch management on a monthly, or more frequent, basis.			
v7	<u>3.4 Deploy Automated Operating System Patch Management Tools</u> Deploy automated software update tools in order to ensure that the operating systems are running the most recent security updates provided by the software vendor.			
v7	<u>3.5 Deploy Automated Software Patch Management Tools</u> Deploy automated software update tools in order to ensure that third-party software on all systems is running the most recent security updates provided by the software vendor.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1195, T1195.001	TA0001	M1051

1.2.2 Ensure gpgcheck is globally activated (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `gpgcheck` option, found in the main section of the `/etc/dnf/dnf.conf` and individual `/etc/yum.repos.d/*` files, determines if an RPM package's signature is checked prior to its installation.

Rationale:

It is important to ensure that an RPM's package signature is always checked prior to installation to ensure that the software is obtained from a trusted source.

Audit:

Global configuration. Run the following command and verify that `gpgcheck` is set to 1:

```
# grep ^gpgcheck /etc/dnf/dnf.conf

gpgcheck=1
```

Configuration in `/etc/yum.repos.d/` takes precedence over the global configuration. Run the following command and verify that there are no instances of entries starting with `gpgcheck` returned set to 0. Nor should there be any invalid (non-boolean) values. When `dnf` encounters such invalid entries they are ignored and the global configuration is applied.

```
# grep -Prs -- '^h*gpgcheck\h*=\h*(0|[2-9]|[1-9][0-9]+|[a-zA-Z_]+)\b'
/etc/yum.repos.d/
```

Remediation:

Edit `/etc/dnf/dnf.conf` and set `gpgcheck=1` in the `[main]` section.

Example:

```
# sed -i 's/^gpgcheck\s*=\s*.*\/gpgcheck=1/' /etc/dnf/dnf.conf
```

Edit any failing files in `/etc/yum.repos.d/*` and set all instances starting with `gpgcheck` to 1.

Example:

```
# find /etc/yum.repos.d/ -name "*.repo" -exec echo "Checking:" {} \; -exec
sed -i 's/^gpgcheck\s*=\s*.*\/gpgcheck=1/' {} \;
```

References:

1. NIST SP 800-53 Rev. 5: SI-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>7.3 Perform Automated Operating System Patch Management</u> Perform operating system updates on enterprise assets through automated patch management on a monthly, or more frequent, basis.			
v7	<u>3.4 Deploy Automated Operating System Patch Management Tools</u> Deploy automated software update tools in order to ensure that the operating systems are running the most recent security updates provided by the software vendor.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1195, T1195.001	TA0005	

1.2.3 Ensure package manager repositories are configured (Manual)

Profile Applicability:

- Level 1 - Server

Description:

Systems need to have the respective package manager repositories configured to ensure that the system is able to receive the latest patches and updates.

Rationale:

If a system's package repositories are misconfigured, important patches may not be identified or a rogue repository could introduce compromised software.

Audit:

Run the following command to verify repositories are configured correctly. The output may vary depending on which repositories are currently configured on the system.

Example:

```
# dnf repolist

repo id                repo name
amazonlinux            Amazon Linux 2023 repository
kernel-livepatch       Amazon Linux 2023 Kernel Livepatch repositor
```

For the repositories in use, inspect the configuration file to ensure all settings are correctly applied according to site policy.

Example:

Depending on the distribution being used the repo file name might differ.

```
cat /etc/yum.repos.d/*.repo
```

Remediation:

Configure your package manager repositories according to site policy.

References:

1. NIST SP 800-53 Rev. 5: SI-2

Additional Information:

For further information about Fedora repositories see: <https://docs.fedoraproject.org/en-US/quick-docs/repositories/>

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>7.3 Perform Automated Operating System Patch Management</u> Perform operating system updates on enterprise assets through automated patch management on a monthly, or more frequent, basis.			
v8	<u>7.4 Perform Automated Application Patch Management</u> Perform application updates on enterprise assets through automated patch management on a monthly, or more frequent, basis.			
v7	<u>3.4 Deploy Automated Operating System Patch Management Tools</u> Deploy automated software update tools in order to ensure that the operating systems are running the most recent security updates provided by the software vendor.			
v7	<u>3.5 Deploy Automated Software Patch Management Tools</u> Deploy automated software update tools in order to ensure that third-party software on all systems is running the most recent security updates provided by the software vendor.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1195, T1195.001	TA0001	M1051

1.2.4 Ensure repo_gpgcheck is globally activated (Manual)

Profile Applicability:

- Level 2 - Server

Description:

The `repo_gpgcheck` option, found in the main section of the `/etc/dnf/dnf.conf` and individual `/etc/yum.repos.d/*` files, will perform a GPG signature check on the repodata.

Rationale:

It is important to ensure that the repository data signature is always checked prior to installation to ensure that the software is not tampered with in any way.

Impact:

Not all repositories, notably RedHat, support `repo_gpgcheck`. Take care to set this value to false (default) for particular repositories that do not support it. If enabled on repositories that do not support `repo_gpgcheck` installation of packages will fail.

Research is required by the user to determine which repositories is configured on the local system and, from that list, which support `repo_gpgcheck`.

Audit:

Global configuration

Run the following command:

```
grep ^repo_gpgcheck /etc/dnf/dnf.conf
```

Verify that `repo_gpgcheck` is set to 1

Per repository configuration

Configuration in `/etc/yum.repos.d/` takes precedence over the global configuration. As an example, to list all the configured repositories, excluding "fedoraproject.org", that specifically disables `repo_gpgcheck`, run the following command:

```
# REPO_URL="fedoraproject.org"
# for repo in $(grep -l "repo_gpgcheck=0" /etc/yum.repos.d/* ); do
#   if ! grep "${REPO_URL}" "${repo}" &>/dev/null; then
#     echo "${repo}"
#   fi
done
```

Per the research that was done on which repositories does not support `repo_gpgcheck`, change the `REPO_URL` variable and run the test.

Remediation:

Global configuration

Edit `/etc/dnf/dnf.conf` and set `repo_gpgcheck=1` in the `[main]` section.

Example:

```
[main]
repo_gpgcheck=1
```

Per repository configuration

First check that the particular repository support GPG checking on the repodata.

Edit any failing files in `/etc/yum.repos.d/*` and set all instances starting with

`repo_gpgcheck` to 1.

References:

1. NIST SP 800-53 Rev. 5: SI-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>7.3 Perform Automated Operating System Patch Management</u> Perform operating system updates on enterprise assets through automated patch management on a monthly, or more frequent, basis.			
v7	<u>3.4 Deploy Automated Operating System Patch Management Tools</u> Deploy automated software update tools in order to ensure that the operating systems are running the most recent security updates provided by the software vendor.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1195, T1195.001	TA0005	

1.3 Filesystem Integrity Checking

AIDE is a file integrity checking tool, similar in nature to Tripwire. While it cannot prevent intrusions, it can detect unauthorized changes to configuration files by alerting when the files are changed. When setting up AIDE, decide internally what the site policy will be concerning integrity checking. Review the AIDE quick start guide and AIDE documentation before proceeding.

1.3.1 Ensure AIDE is installed (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Advanced Intrusion Detection Environment (AIDE) is a intrusion detection tool that uses predefined rules to check the integrity of files and directories in the Linux operating system. AIDE has its own database to check the integrity of files and directories.

AIDE takes a snapshot of files and directories including modification times, permissions, and file hashes which can then be used to compare against the current state of the filesystem to detect modifications to the system.

Rationale:

By monitoring the filesystem state compromised files can be detected to prevent or limit the exposure of accidental or malicious misconfigurations or modified binaries.

Audit:

Run the following command and verify `aide` is installed:

```
# rpm -q aide  
  
aide-<version>
```

Remediation:

Run the following command to install AIDE:

```
# dnf install aide
```

Configure AIDE as appropriate for your environment. Consult the AIDE documentation for options.

Initialize AIDE:

Run the following commands:

```
# aide --init  
# mv /var/lib/aide/aide.db.new.gz /var/lib/aide/aide.db.gz
```

References:

1. AIDE stable manual: <http://aide.sourceforge.net/stable/manual.html>
2. NIST SP 800-53 Rev. 5: AU-2

Additional Information:

The prelinking feature can interfere with AIDE because it alters binaries to speed up their start up times. Run `prelink -ua` to restore the binaries to their prelinked state, thus avoiding false positives from AIDE.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.14 <u>Log Sensitive Data Access</u> Log sensitive data access, including modification and disposal.			●
v7	14.9 <u>Enforce Detail Logging for Access or Changes to Sensitive Data</u> Enforce detailed audit logging for access to sensitive data or changes to sensitive data (utilizing tools such as File Integrity Monitoring or Security Information and Event Monitoring).			●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1565, T1565.001	TA0001	M1022

1.3.2 Ensure filesystem integrity is regularly checked (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Periodic checking of the filesystem integrity is needed to detect changes to the filesystem.

Rationale:

Periodic file checking allows the system administrator to determine on a regular basis if critical files have been changed in an unauthorized fashion.

Audit:

Run the following commands to verify a cron job scheduled to run the aide check.

```
# grep -Ers '^(#[^#]+\s+)?(\/usr\/s?bin\/|^\/s*)aide(\/.wrapper)?\s(--?\S+\s)*(-
-(check|update)|\${AIDEARGS})\b' /etc/cron.* /etc/crontab /var/spool/cron/
```

Ensure a cron job in compliance with site policy is returned.

OR run the following commands to verify that aidecheck.service and aidecheck.timer are enabled and aidcheck.timer is running

```
# systemctl is-enabled aidecheck.service
# systemctl is-enabled aidecheck.timer
# systemctl status aidecheck.timer
```

Remediation:

If cron will be used to schedule and run aide check

Run the following command:

```
# crontab -u root -e
```

Add the following line to the crontab:

```
0 5 * * * /usr/sbin/aide --check
```

OR if aidecheck.service and aidecheck.timer will be used to schedule and run aide check:

Create or edit the file `/etc/systemd/system/aidecheck.service` and add the following lines:

```
[Unit]
Description=Aide Check

[Service]
Type=simple
ExecStart=/usr/sbin/aide --check

[Install]
WantedBy=multi-user.target
```

Create or edit the file `/etc/systemd/system/aidecheck.timer` and add the following lines:

```
[Unit]
Description=Aide check every day at 5AM

[Timer]
OnCalendar=*-*-* 05:00:00
Unit=aidecheck.service

[Install]
WantedBy=multi-user.target
```

Run the following commands:

```
# chown root:root /etc/systemd/system/aidecheck.*
# chmod 0644 /etc/systemd/system/aidecheck.*

# systemctl daemon-reload

# systemctl enable aidecheck.service
# systemctl --now enable aidecheck.timer
```

References:

1. <https://github.com/konstruktoid/hardening/blob/master/config/aidecheck.service>
2. <https://github.com/konstruktoid/hardening/blob/master/config/aidecheck.timer>
3. NIST SP 800-53 Rev. 5: AU-2

Additional Information:

The checking in this recommendation occurs every day at 5am. Alter the frequency and time of the checks in compliance with site policy.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.14 Log Sensitive Data Access Log sensitive data access, including modification and disposal.			●
v7	14.9 Enforce Detail Logging for Access or Changes to Sensitive Data Enforce detailed audit logging for access to sensitive data or changes to sensitive data (utilizing tools such as File Integrity Monitoring or Security Information and Event Monitoring).			●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1036, T1036.005	TA0040	M1022

1.3.3 Ensure cryptographic mechanisms are used to protect the integrity of audit tools (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Audit tools include, but are not limited to, vendor-provided and open source audit tools needed to successfully view and manipulate audit information system activity and records. Audit tools include custom queries and report generators.

Rationale:

Protecting the integrity of the tools used for auditing purposes is a critical step toward ensuring the integrity of audit information. Audit information includes all information (e.g., audit records, audit settings, and audit reports) needed to successfully audit information system activity.

Attackers may replace the audit tools or inject code into the existing tools with the purpose of providing the capability to hide or erase system activity from the audit logs.

Audit tools should be cryptographically signed in order to provide the capability to identify when the audit tools have been modified, manipulated, or replaced. An example is a checksum hash of the file or files.

Audit:

Verify that Advanced Intrusion Detection Environment (AIDE) is properly configured . Run the following command to verify that AIDE is configured to use cryptographic mechanisms to protect the integrity of audit tools:

```
# grep -Ps -- '(\sbin\/(audit|au)\H*\b)' /etc/aide.conf.d/*.conf  
/etc/aide.conf
```

Verify the output includes:

```
/sbin/auditctl p+i+n+u+g+s+b+acl+xattrs+sha512  
/sbin/auditd p+i+n+u+g+s+b+acl+xattrs+sha512  
/sbin/ausearch p+i+n+u+g+s+b+acl+xattrs+sha512  
/sbin/aureport p+i+n+u+g+s+b+acl+xattrs+sha512  
/sbin/autrace p+i+n+u+g+s+b+acl+xattrs+sha512  
/sbin/auditrules p+i+n+u+g+s+b+acl+xattrs+sha512
```

Remediation:

Add or update the following selection lines for to a file ending in `.conf` in the `/etc/aide.conf.d/` directory or to `/etc/aide.conf` to protect the integrity of the audit tools:

```
# Audit Tools
/sbin/auditctl p+i+n+u+g+s+b+acl+xattrs+sha512
/sbin/auditd p+i+n+u+g+s+b+acl+xattrs+sha512
/sbin/ausearch p+i+n+u+g+s+b+acl+xattrs+sha512
/sbin/aureport p+i+n+u+g+s+b+acl+xattrs+sha512
/sbin/autrace p+i+n+u+g+s+b+acl+xattrs+sha512
/sbin/augenrules p+i+n+u+g+s+b+acl+xattrs+sha512
```

References:

1. NIST SP 800-53 Rev. 5: AU-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.14 <u>Log Sensitive Data Access</u> Log sensitive data access, including modification and disposal.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1083, T1083.000	TA0007	M1022

1.4 Secure Boot Settings

The recommendations in this section focus on securing the bootloader and settings involved in the boot process directly.

Note:

- In Fedora 28 based distributions, the kernel command-line parameters for systems using the GRUB2 bootloader were defined in the `kernelopts` environment variable. This variable was stored in the `/boot/grub2/grubenv` file for each kernel boot entry. However, storing the kernel command-line parameters using `kernelopts` was not robust. Therefore, the `kernelopts` has been removed and the kernel command-line parameters are now stored in the Boot Loader Specification (BLS) snippet, instead of in the `/boot/loader/entries/<KERNEL_BOOT_ENTRY>.conf` file.
- Boot loader configuration files are unified across CPU architectures
 - Configuration files for the GRUB boot loader are now stored in the `/boot/grub2/` directory on all supported CPU architectures. The `/boot/efi/EFI/redhat/grub.cfg` file, which GRUB previously used as the main configuration file on UEFI systems, now simply loads the `/boot/grub2/grub.cfg` file.
 - This change simplifies the layout of the GRUB configuration file, improves user experience, and provides the following notable benefits:
 - You can boot the same installation with either EFI or legacy BIOS.
 - You can use the same documentation and commands for all architectures.
 - GRUB configuration tools are more robust, because they no longer rely on symbolic links and they do not have to handle platform-specific cases.
 - The usage of the GRUB configuration files is aligned with images generated by CoreOS Assembler (COSA) and OSBuild.
 - The usage of the GRUB configuration files is aligned with other Linux distributions.
 - Fedora 34 based distributions no longer boot on 32-bit UEFI
- Support for the 32-bit UEFI firmware was removed from the GRUB and shim boot loaders. As a consequence, Fedora 34 based distributions requires a 64-bit UEFI, and can no longer boot on 64-bit systems that use a 32-bit UEFI.
 - The following packages have been removed as part of this change:
 - `grub2-efi-ia32`
 - `grub2-efi-ia32-cdboot`
 - `grub2-efi-ia32-modules`
 - `shim-ia32`

Reference: https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html/considerations_in_adopting_rhel_9/assembly_kernel_considerations-in-adopting-rhel-9

1.4.1 Ensure permissions on bootloader config are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The grub files contain information on boot settings and passwords for unlocking boot options.

Rationale:

Setting the permissions to read and write for root only prevents non-root users from seeing the boot parameters or changing them. Non-root users who read the boot parameters may be able to identify weaknesses in security upon boot and be able to exploit them.

Audit:

Run the following script to verify grub configuration files:

- For systems using UEFI (Files located in `/boot/efi/EFI/*`):
 - Mode is `0700` or more restrictive
- For systems using BIOS (Files located in `/boot/grub2/*`):
 - Mode is `0600` or more restrictive
- Owner is the user `root`
- Group owner is group `root`

```

#!/usr/bin/env bash

{
    l_output="" l_output2=""
    file_mug_chk()
    {
        l_out="" l_out2=""
        [[ "$(dirname "$l_file")" =~ ^\./boot\./efi\./EFI ]] && l_pmask="0077" ||
l_pmask="0177"
        l_maxperm="$( printf '%o' $(( 0777 & ~$l_pmask )) )"
        if [ $(( $l_mode & $l_pmask )) -gt 0 ]; then
            l_out2="$l_out2\n    - Is mode \"$l_mode\" and should be mode:
\"$l_maxperm\" or more restrictive"
        else
            l_out="$l_out\n    - Is correctly mode: \"$l_mode\" which is mode:
\"$l_maxperm\" or more restrictive"
        fi
        if [ "$l_user" = "root" ]; then
            l_out="$l_out\n    - Is correctly owned by user: \"$l_user\""
        else
            l_out2="$l_out2\n    - Is owned by user: \"$l_user\" and should be
owned by user: \"root\""
        fi
        if [ "$l_group" = "root" ]; then
            l_out="$l_out\n    - Is correctly group-owned by group: \"$l_user\""
        else
            l_out2="$l_out2\n    - Is group-owned by group: \"$l_user\" and
should be group-owned by group: \"root\""
        fi
        [ -n "$l_out" ] && l_output="$l_output\n    - File: \"$l_file\"$l_out\n"
        [ -n "$l_out2" ] && l_output2="$l_output2\n    - File:
\"$l_file\"$l_out2\n"
    }
    while IFS= read -r -d $'\0' l_gfile; do
        while read -r l_file l_mode l_user l_group; do
            file_mug_chk
            done <<< "$(stat -Lc '%n %a %U %G' "$l_gfile")"
            done <<(find /boot -type f \( -name 'grub*' -o -name 'user.cfg' \) -
print0)
            if [ -z "$l_output2" ]; then
                echo -e "\n- Audit Result:\n    *** PASS ***\n- * Correctly set *
:\n$l_output\n"
            else
                echo -e "\n- Audit Result:\n    ** FAIL **\n- * Reasons for audit
failure * :\n$l_output2\n"
                [ -n "$l_output" ] && echo -e " - * Correctly set * :\n$l_output\n"
            fi
        done
    done
}

```

Remediation:

Run the following to update the mode, ownership, and group ownership of the grub configuration files:

-- IF -- the system uses UEFI (Files located in /boot/efi/EFI/*)

Edit /etc/fstab and add the `umask=0077`, `uid=0`, and `gid=0` options:

Example:

```
<device> /boot/efi vfat defaults,umask=0027,fmask=0077,uid=0,gid=0 0 0
```

Note: This may require a re-boot to enable the change

-- OR --

-- IF -- the system uses BIOS (Files located in /boot/grub2/*)

Run the following commands to set ownership and permissions on your grub configuration file(s):

```
# [ -f /boot/grub2/grub.cfg ] && chown root:root /boot/grub2/grub.cfg
# [ -f /boot/grub2/grub.cfg ] && chmod u-x,go-rwx /boot/grub2/grub.cfg

# [ -f /boot/grub2/grubenv ] && chown root:root /boot/grub2/grubenv
# [ -f /boot/grub2/grubenv ] && chmod u-x,go-rwx /boot/grub2/grubenv

# [ -f /boot/grub2/user.cfg ] && chown root:root /boot/grub2/user.cfg
# [ -f /boot/grub2/user.cfg ] && chmod u-x,go-rwx /boot/grub2/user.cfg
```

Additional Information:

This recommendation is designed around the grub bootloader, if LILO or another bootloader is in use in your environment enact equivalent settings.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 Configure Data Access Control Lists Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 Protect Information through Access Control Lists Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1542, T1542.000	TA0005, TA0007	M1022

1.5 Additional Process Hardening

1.5.1 Ensure address space layout randomization (ASLR) is enabled (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Address space layout randomization (ASLR) is an exploit mitigation technique which randomly arranges the address space of key data areas of a process.

Rationale:

Randomly placing virtual memory regions will make it difficult to write memory page exploits as the memory placement will be consistently shifting.

Audit:

Run the following script to verify the following kernel parameter is set in the running configuration and correctly loaded from a kernel parameter configuration file:

- `kernel.randomize_va_space` is set to 2

Note: kernel parameters are loaded by file and parameter order precedence. The following script observes this precedence as part of the auditing procedure. The parameters being checked may be set correctly in a file. If that file is superseded, the parameter is overridden by an incorrect setting later in that file, or in a conically later file, that "correct" setting will be ignored both by the script and by the system during a normal kernel parameter load sequence.

```

#!/usr/bin/env bash

{
    l_output="" l_output2=""
    a_parlist=("kernel.randomize_va_space=2")
    l_ufwscf="$([ -f /etc/default/ufw ] && awk -F= '/^\s*IPT_SYSCTL=/ {print $2}'
/etc/default/ufw)"
    kernel_parameter_chk()
    {
        l_krp="$(sysctl "$l_kpname" | awk -F= '{print $2}' | xargs)" # Check running configuration
        if [ "$l_krp" = "$l_kpvalue" ]; then
            l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_krp\" in the running
configuration"
        else
            l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_krp\" in the running
configuration and should have a value of: \"$l_kpvalue\""
        fi
        unset A_out; declare -A A_out # Check durable setting (files)
        while read -r l_out; do
            if [ -n "$l_out" ]; then
                if [[ $l_out =~ ^\s*# ]]; then
                    l_file="$l_out// # /)"
                else
                    l_kpar="$(awk -F= '{print $1}' <<< "$l_out" | xargs)"
                    [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]=" $l_file")
                fi
            fi
            done < <(/usr/lib/systemd/systemd-sysctl --cat-config | grep -Po
'^\h*([^\#\n\r]+|#\h*\[/^\#\n\r\h]+\.\conf\b)')
            if [ -n "$l_ufwscf" ]; then # Account for systems with UFW (Not covered by systemd-sysctl -
-cat-config)
                l_kpar="$(grep -Po '^ \h*$l_kpname\b' "$l_ufwscf" | xargs)"
                l_kpar="$l_kpar//\./)"
                [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]=" $l_ufwscf")
            fi
            if (( ${#A_out[@]} > 0 )); then # Assess output from files and generate output
                while IFS="" read -r l_fkpname l_fkpvalue; do
                    l_fkpname="$l_fkpname// /)"; l_fkpvalue="$l_fkpvalue// /)"
                    if [ "$l_fkpvalue" = "$l_kpvalue" ]; then
                        l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_krp\" in \"$(printf
'%s' "${A_out[@]}")\""\n"
                    else
                        l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_fkpvalue\" in
\"$(printf '%s' "${A_out[@]}")\" and should have a value of: \"$l_kpvalue\""\n"
                    fi
                    done < <(grep -Po -- "\h*$l_kpname\h*=\h*\H+" "${A_out[@]}")
                else
                    l_output2="$l_output2\n - \"$l_kpname\" is not set in an included file\n    ** Note:
\"$l_kpname\" May be set in a file that's ignored by load procedure **\n"
                fi
            }
            while IFS="" read -r l_kpname l_kpvalue; do # Assess and check parameters
                l_kpname="$l_kpname// /)"; l_kpvalue="$l_kpvalue// /)"
                if ! grep -Pqs '\h*0\b' /sys/module/ipv6/parameters/disable && grep -q '^net.ipv6.' <<<
"$l_kpname"; then
                    l_output="$l_output\n - IPv6 is disabled on the system, \"$l_kpname\" is not applicable"
                else
                    kernel_parameter_chk
                fi
            done < <(printf '%s\n' "${a_parlist[@]}")
            if [ -z "$l_output2" ]; then # Provide output from checks
                echo -e "\n- Audit Result:\n    ** PASS **\n$l_output\n"
            else
                echo -e "\n- Audit Result:\n    ** FAIL **\n - Reason(s) for audit failure:\n$l_output2\n"
                [ -n "$l_output" ] && echo -e "\n- Correctly set:\n$l_output\n"
            fi
        }
    }
}

```

Remediation:

Set the following parameter in `/etc/sysctl.conf` or a file in `/etc/sysctl.d/` ending in `.conf`:

- `kernel.randomize_va_space = 2`

Example:

```
# printf "  
kernel.randomize_va_space = 2  
" >> /etc/sysctl.d/60-kernel_sysctl.conf
```

Run the following command to set the active kernel parameter:

```
# sysctl -w kernel.randomize_va_space=2
```

Note: If these settings appear in a conically later file, or later in the same file, these settings will be overwritten

Default Value:

`kernel.randomize_va_space = 2`

References:

- CCI-000366: The organization implements the security configuration settings
- NIST SP 800-53 Rev. 5: CM-6

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	10.5 Enable Anti-Exploitation Features Enable anti-exploitation features on enterprise assets and software, where possible, such as Microsoft® Data Execution Prevention (DEP), Windows® Defender Exploit Guard (WDEG), or Apple® System Integrity Protection (SIP) and Gatekeeper™.		●	●
v7	8.3 Enable Operating System Anti-Exploitation Features/ Deploy Anti-Exploit Technologies Enable anti-exploitation features such as Data Execution Prevention (DEP) or Address Space Layout Randomization (ASLR) that are available in an operating system or deploy appropriate toolkits that can be configured to apply protection to a broader set of applications and executables.		●	●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1068, T1068.000	TA0002	M1050

1.5.2 Ensure `ptrace_scope` is restricted (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `ptrace()` system call provides a means by which one process (the "tracer") may observe and control the execution of another process (the "tracee"), and examine and change the tracee's memory and registers.

Rationale:

If one application is compromised, it would be possible for an attacker to attach to other running processes (e.g. Bash, Firefox, SSH sessions, GPG agent, etc) to extract additional credentials and continue to expand the scope of their attack.

Enabling restricted mode will limit the ability of a compromised process to `PTRACE_ATTACH` on other processes running under the same user. With restricted mode, `ptrace` will continue to work with root user.

Audit:

Run the following script to verify the following kernel parameter is set in the running configuration and correctly loaded from a kernel parameter configuration file:

- `kernel.yama.ptrace_scope` is set to 1

Note: kernel parameters are loaded by file and parameter order precedence. The following script observes this precedence as part of the auditing procedure. The parameters being checked may be set correctly in a file. If that file is superseded, the parameter is overridden by an incorrect setting later in that file, or in a conically later file, that "correct" setting will be ignored both by the script and by the system during a normal kernel parameter load sequence.

```
#!/usr/bin/env bash

{
    l_output="" l_output2=""
    a_parlist=("kernel.yama.pttrace_scope=1")
    l_ufwscf="$([ -f /etc/default/ufw ] && awk -F= '/^\s*IPT_SYSCTL=/ {print $2}'
/etc/default/ufw)"
    kernel_parameter_chk()
    {
        l_krp="$(sysctl "$l_kpname" | awk -F= '{print $2}' | xargs)" # Check running configuration
        if [ "$l_krp" = "$l_kpvalue" ]; then
            l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_krp\" in the running
configuration"
        else
            l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_krp\" in the running
configuration and should have a value of: \"$l_kpvalue\""
        fi
        unset A_out; declare -A A_out # Check durable setting (files)
        while read -r l_out; do
            if [ -n "$l_out" ]; then
                if [[ $l_out =~ ^\s*# ]]; then
                    l_file="$l_out// # /)"
                else
                    l_kpar="$(awk -F= '{print $1}' <<< "$l_out" | xargs)"
                    [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]=" $l_file")
                fi
            fi
            done < <(/usr/lib/systemd/systemd-sysctl --cat-config | grep -Po
'^\h*([^\#\n\r]+|#\h*\[/^\#\n\r\h]+\.\conf\b)')
            if [ -n "$l_ufwscf" ]; then # Account for systems with UFW (Not covered by systemd-sysctl -
-cat-config)
                l_kpar="$(grep -Po '^ \h*$l_kpname\b' "$l_ufwscf" | xargs)"
                l_kpar="$l_kpar//\./)"
                [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]=" $l_ufwscf")
            fi
            if (( ${#A_out[@]} > 0 )); then # Assess output from files and generate output
                while IFS="" read -r l_fkpname l_fkpvalue; do
                    l_fkpname="$l_fkpname// /)"; l_fkpvalue="$l_fkpvalue// /)"
                    if [ "$l_fkpvalue" = "$l_kpvalue" ]; then
                        l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_krp\" in \"$(printf
'%s' "${A_out[@]}")\""\n"
                    else
                        l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_fkpvalue\" in
\"$(printf '%s' "${A_out[@]}")\" and should have a value of: \"$l_kpvalue\""\n"
                    fi
                    done < <(grep -Po -- "\^ \h*$l_kpname\h*=\h*\H+" "${A_out[@]}")
                else
                    l_output2="$l_output2\n - \"$l_kpname\" is not set in an included file\n    ** Note:
\"$l_kpname\" May be set in a file that's ignored by load procedure **\n"
                fi
            }
            while IFS="" read -r l_kpname l_kpvalue; do # Assess and check parameters
                l_kpname="$l_kpname// /)"; l_kpvalue="$l_kpvalue// /)"
                if ! grep -Pqs '\^ \h*0\b' /sys/module/ipv6/parameters/disable && grep -q '\^net.ipv6.' <<<
"$l_kpname"; then
                    l_output="$l_output\n - IPv6 is disabled on the system, \"$l_kpname\" is not applicable"
                else
                    kernel_parameter_chk
                fi
            done < <(printf '%s\n' "${a_parlist[@]}")
            if [ -z "$l_output2" ]; then # Provide output from checks
                echo -e "\n- Audit Result:\n    ** PASS **\n$l_output\n"
            else
                echo -e "\n- Audit Result:\n    ** FAIL **\n - Reason(s) for audit failure:\n$l_output2\n"
                [ -n "$l_output" ] && echo -e "\n- Correctly set:\n$l_output\n"
            fi
        }
    }
}
```

Remediation:

Set the following parameter in `/etc/sysctl.conf` or a file in `/etc/sysctl.d/` ending in `.conf`:

- `kernel.yama.ptrace_scope = 1`

Example:

```
# printf "  
kernel.yama.ptrace_scope = 1  
" >> /etc/sysctl.d/60-kernel_sysctl.conf
```

Run the following command to set the active kernel parameter:

```
# sysctl -w kernel.yama.ptrace_scope=1
```

Note: If these settings appear in a conically later file, or later in the same file, these settings will be overwritten

Default Value:

`kernel.yama.ptrace_scope=1`

References:

1. <https://www.kernel.org/doc/Documentation/security/Yama.txt>
2. <https://github.com/raj3shp/termspy>
3. NIST SP 800-53 Rev. 5: CM-6

Additional Information:

Ptrace is very rarely used by regular applications and is mostly used by debuggers such as `gdb` and `strace`.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 <u>Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●
v7	9.2 <u>Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.		●	●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1055.008	TA0005	M1040

1.5.3 Ensure core dump storage is disabled (Automated)

Profile Applicability:

- Level 1 - Server

Description:

A core dump is the memory of an executable program. It is generally used to determine why a program aborted. It can also be used to glean confidential information from a core file.

Rationale:

A core dump includes a memory image taken at the time the operating system terminates an application. The memory image could contain sensitive data and is generally useful only for developers trying to debug problems.

Audit:

Run the following command to verify `Storage` is set to `none` in `/etc/systemd/coredump.conf`:

```
# grep -i '^s*storage\s*=\s*none' /etc/systemd/coredump.conf
Storage=none
```

Remediation:

Edit `/etc/systemd/coredump.conf` and edit or add the following line:

```
Storage=none
```

References:

1. <https://www.freedesktop.org/software/systemd/man/coredump.conf.html>
2. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v7	13.2 <u>Remove Sensitive Data or Systems Not Regularly Accessed by Organization</u> Remove sensitive data or systems not regularly accessed by the organization from the network. These systems shall only be used as stand alone systems (disconnected from the network) by the business unit needing to occasionally use the system or completely virtualized and powered off until needed.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1005, T1005.000	TA0007	M1057

1.5.4 Ensure core dump backtraces are disabled (Automated)

Profile Applicability:

- Level 1 - Server

Description:

A core dump is the memory of an executable program. It is generally used to determine why a program aborted. It can also be used to glean confidential information from a core file.

Rationale:

A core dump includes a memory image taken at the time the operating system terminates an application. The memory image could contain sensitive data and is generally useful only for developers trying to debug problems, increasing the risk to the system.

Audit:

Run the following command to verify `ProcessSizeMax` is set to 0 in

`/etc/systemd/coredump.conf`:

```
# grep -Pi '^h*ProcessSizeMax\h*=\h*0\b' /etc/systemd/coredump.conf || echo
-e "\n- Audit results:\n  ** Fail **\n - \"ProcessSizeMax\" is: \"$(grep -i
'ProcessSizeMax' /etc/systemd/coredump.conf)\""
```

Remediation:

Edit or add the following line in `/etc/systemd/coredump.conf`:

```
ProcessSizeMax=0
```

Default Value:

`ProcessSizeMax=2G`

References:

1. <https://www.freedesktop.org/software/systemd/man/coredump.conf.html>
2. NIST SP 800-53 Rev. 5: CM-6

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v7	13.2 <u>Remove Sensitive Data or Systems Not Regularly Accessed by Organization</u> Remove sensitive data or systems not regularly accessed by the organization from the network. These systems shall only be used as stand alone systems (disconnected from the network) by the business unit needing to occasionally use the system or completely virtualized and powered off until needed.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1005, T1005.000	TA0007	M1057

1.6 Mandatory Access Control

Mandatory Access Control (MAC) provides an additional layer of access restrictions to processes on top of the base Discretionary Access Controls. By restricting how processes can access files and resources on a system the potential impact from vulnerabilities in the processes can be reduced.

Impact: Mandatory Access Control limits the capabilities of applications and daemons on a system, while this can prevent unauthorized access the configuration of MAC can be complex and difficult to implement correctly preventing legitimate access from occurring.

1.6.1 Configure SELinux

SELinux implements Mandatory Access Control (MAC). Every process and system resource has a special security label called an SELinux context. A SELinux context, sometimes referred to as an SELinux label, is an identifier which abstracts away the system-level details and focuses on the security properties of the entity. Not only does this provide a consistent way of referencing objects in the SELinux policy, but it also removes any ambiguity that can be found in other identification methods. For example, a file can have multiple valid path names on a system that makes use of bind mounts.

The SELinux policy uses these contexts in a series of rules which define how processes can interact with each other and the various system resources. By default, the policy does not allow any interaction unless a rule explicitly grants access.

In Fedora 34 Family Linux distributions, system services are controlled by the `systemd` daemon; `systemd` starts and stops all services, and users and processes communicate with `systemd` using the `systemctl` utility. The `systemd` daemon can consult the SELinux policy and check the label of the calling process and the label of the unit file that the caller tries to manage, and then ask SELinux whether or not the caller is allowed the access. This approach strengthens access control to critical system capabilities, which include starting and stopping system services.

This automatically limits the damage that the software can do to files accessible by the calling user. The user does not need to take any action to gain this benefit. For an action to occur, both the traditional DAC permissions must be satisfied as well as the SELinux MAC rules. The action will not be allowed if either one of these models does not permit the action. In this way, SELinux rules can only make a system's permissions more restrictive and secure. SELinux requires a complex policy to allow all the actions required of a system under normal operation. Two such policies have been designed for use with Fedora 34 Family Linux distributions and are included with the system:

`targeted` and `mls`. These are described as follows:

- `targeted`: `targeted` processes run in their own domain, called a confined domain. In a confined domain, the files that a `targeted` process has access to are limited. If a confined process is compromised by an attacker, the attacker's access to resources and the possible damage they can do is also limited. SELinux denies access to these resources and logs the denial.
- `mls`: implements Multi-Level Security (MLS), which introduces even more kinds of labels (sensitivity and category) and rules that govern access based on these.

This section provides guidance for the configuration of the `targeted` policy.

Notes:

- Remember that SELinux policy rules are checked after DAC rules. SELinux policy rules are not used if DAC rules deny access first, which means that no SELinux denial is logged if the traditional DAC rules prevent the access.
- This section only applies if SELinux is in use on the system. Additional Mandatory Access Control systems exist.
- To avoid incorrect SELinux labeling and subsequent problems, ensure that you start services using a `systemctl start` command.

References:

1. NSA SELinux resources:
 1. <http://www.nsa.gov/research/selinux>
 2. <http://www.nsa.gov/research/selinux/list.shtml>
2. Fedora SELinux resources:
 1. FAQ: <http://docs.fedoraproject.org/selinux-faq>
 2. User Guide: https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html-single/using_selinux/index
 3. Managing Services Guide: <http://docs.fedoraproject.org/selinux-managing-confined-services-guide>
3. SELinux Project web page and wiki:
 1. <http://www.selinuxproject.org>

1.6.1.1 Ensure SELinux is installed (Automated)

Profile Applicability:

- Level 1 - Server

Description:

SELinux provides Mandatory Access Control.

Rationale:

Without a Mandatory Access Control system installed only the default Discretionary Access Control system will be available.

Audit:

Verify SELinux is installed.
Run the following command:

```
# rpm -q libselinux  
  
libselinux-<version>
```

Remediation:

Run the following command to install SELinux:

```
# dnf install libselinux
```

References:

1. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1068, T1068.000	TA0003	M1026

1.6.1.2 Ensure SELinux is not disabled in bootloader configuration (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Configure SELINUX to be enabled at boot time and verify that it has not been overwritten by the grub boot parameters.

Rationale:

SELinux must be enabled at boot time in your grub configuration to ensure that the controls it provides are not overridden.

Impact:

Files created while SELinux is disabled are not labeled at all. This behavior causes problems when changing to enforcing mode because files are labeled incorrectly or are not labeled at all. To prevent incorrectly labeled and unlabeled files from causing problems, file systems are automatically relabeled when changing from the disabled state to permissive or enforcing mode. This can be a long running process that should be accounted for as it may extend downtime during initial re-boot.

Audit:

Run the following command to verify that neither the `selinux=0` or `enforcing=0` parameters have been set:

```
# grubby --info=ALL | grep -Po '(selinux|enforcing)=0\b'
```

Nothing should be returned

Remediation:

Run the following command to remove the `selinux=0` and `enforcing=0` parameters:

```
grubby --update-kernel ALL --remove-args "selinux=0 enforcing=0"
```

Run the following command to remove the `selinux=0` and `enforcing=0` parameters if they were created by the deprecated `grub2-mkconfig` command:

```
# grep -Prsq --
'\h*([\#\n\r]+\h+)?kernelopts=([\#\n\r]+\h+)?(selinux|enforcing)=0\b'
/boot/grub2 /boot/efi && grub2-mkconfig -o "$(grep -Pr1 --
'\h*([\#\n\r]+\h+)?kernelopts=([\#\n\r]+\h+)?(selinux|enforcing)=0\b'
/boot/grub2 /boot/efi)"
```

References:

1. NIST SP 800-53 Rev. 5: SI-7

Additional Information:

This recommendation is designed around the grub 2 bootloader, if another bootloader is in use in your environment enact equivalent settings.

`grubby` is a command line tool for updating and displaying information about the configuration files for the grub2 and zipl boot loaders. It is primarily designed to be used from scripts which install new kernels and need to find information about the current boot environment.

- All bootloaders define the boot entries as individual configuration fragments that are stored by default in `/boot/loader/entries`. The format for the config files is specified at https://systemd.io/BOOT_LOADER_SPECIFICATION. The `grubby` tool is used to update and display the configuration defined in the `BootLoaderSpec` fragment files.
- There are a number of ways to specify the kernel used for `--info`, `--remove-kernel`, and `--update-kernel`. Specifying `DEFAULT` or `ALL` selects the de-fault entry and all of the entries, respectively. Also, the title of a boot entry may be specified by using `TITLE=title` as the argument; all entries with that title are used.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1068, T1068.000	TA0003	M1026

1.6.1.3 Ensure SELinux policy is configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Configure SELinux to meet or exceed the default targeted policy, which constrains daemons and system software only.

Rationale:

Security configuration requirements vary from site to site. Some sites may mandate a policy that is stricter than the default policy, which is perfectly acceptable. This item is intended to ensure that at least the default recommendations are met.

Audit:

Run the following commands and ensure output matches either "targeted" or "mls":

```
# grep -E '^s*SELINUXTYPE=(targeted|mls)\b' /etc/selinux/config
SELINUXTYPE=targeted
# sestatus | grep Loaded
Loaded policy name:                targeted
```

Remediation:

Edit the `/etc/selinux/config` file to set the SELINUXTYPE parameter:

```
SELINUXTYPE=targeted
```

References:

1. NIST SP 800-53 Rev. 5: AC-3, MP-2

Additional Information:

If your organization requires stricter policies, ensure that they are set in the `/etc/selinux/config` file.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1068, T1068.000	TA0005	M1048

1.6.1.4 Ensure the SELinux mode is not disabled (Automated)

Profile Applicability:

- Level 1 - Server

Description:

SELinux can run in one of three modes: disabled, permissive, or enforcing:

- Enforcing - Is the default, and recommended, mode of operation; in enforcing mode SELinux operates normally, enforcing the loaded security policy on the entire system.
- Permissive - The system acts as if SELinux is enforcing the loaded security policy, including labeling objects and emitting access denial entries in the logs, but it does not actually deny any operations. While not recommended for production systems, permissive mode can be helpful for SELinux policy development.
- Disabled - Is strongly discouraged; not only does the system avoid enforcing the SELinux policy, it also avoids labeling any persistent objects such as files, making it difficult to enable SELinux in the future

Note: you can set individual domains to permissive mode while the system runs in enforcing mode. For example, to make the httpd_t domain permissive:

```
# semanage permissive -a httpd_t
```

Rationale:

Running SELinux in disabled mode is strongly discouraged; not only does the system avoid enforcing the SELinux policy, it also avoids labeling any persistent objects such as files, making it difficult to enable SELinux in the future.

Audit:

Run the following commands and ensure output matches:
Run the following command to verify SELinux's current mode:

```
# getenforce  
  
Enforcing  
-OR-  
Permissive
```

Run the following command to verify SELinux's configured mode:

```
# grep -Ei '^s*SELINUX=(enforcing|permissive)' /etc/selinux/config  
  
SELINUX=enforcing  
-OR-  
SELINUX=permissive
```

Remediation:

Run one of the following commands to set SELinux's running mode:
To set SELinux mode to `Enforcing`:

```
# setenforce 1
```

OR

To set SELinux mode to `Permissive`:

```
# setenforce 0
```

Edit the `/etc/selinux/config` file to set the SELINUX parameter:
For `Enforcing` mode:

```
SELINUX=enforcing
```

OR

For `Permissive` mode:

```
SELINUX=permissive
```

References:

1. https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/selinux_users_and_administrators_guide/sect-5-4-security-enhanced_linux-introduction-selinux_modes
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1068, T1068.000, T1565, T1565.001, T1565.003	TA0003	M1026

1.6.1.5 Ensure the SELinux mode is enforcing (Automated)

Profile Applicability:

- Level 2 - Server

Description:

SELinux can run in one of three modes: disabled, permissive, or enforcing:

- Enforcing - Is the default, and recommended, mode of operation; in enforcing mode SELinux operates normally, enforcing the loaded security policy on the entire system.
- Permissive - The system acts as if SELinux is enforcing the loaded security policy, including labeling objects and emitting access denial entries in the logs, but it does not actually deny any operations. While not recommended for production systems, permissive mode can be helpful for SELinux policy development.
- Disabled - Is strongly discouraged; not only does the system avoid enforcing the SELinux policy, it also avoids labeling any persistent objects such as files, making it difficult to enable SELinux in the future

Note: you can set individual domains to permissive mode while the system runs in enforcing mode. For example, to make the httpd_t domain permissive:

```
# semanage permissive -a httpd_t
```

Rationale:

Running SELinux in disabled mode the system not only avoids enforcing the SELinux policy, it also avoids labeling any persistent objects such as files, making it difficult to enable SELinux in the future.

Running SELinux in Permissive mode, though helpful for developing SELinux policy, only logs access denial entries, but does not deny any operations.

Audit:

Run the following commands and ensure output matches:
Run the following command to verify SELinux's current mode:

```
# getenforce
Enforcing
```

Run the following command to verify SELinux's configured mode:

```
# grep -i SELINUX=enforcing /etc/selinux/config
SELINUX=enforcing
```

Remediation:

Run the following command to set SELinux's running mode:

```
# setenforce 1
```

Edit the `/etc/selinux/config` file to set the SELINUX parameter:
For Enforcing mode:

```
SELINUX=enforcing
```

References:

1. https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/selinux_users_and_administrators_guide/sect-3.3-configure-data-access-control-lists
2. CCI-002165: The information system enforces organization-defined discretionary access control policies over defined subjects and objects.
3. NIST SP 800-53 Rev. 5: AC-3, SI-6
4. CCI-002696: The information system verifies correct operation of organization-defined security functions.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 Configure Data Access Control Lists Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 Protect Information through Access Control Lists Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1068, T1068.000, T1565, T1565.001, T1565.003	TA0005	M1026

1.6.1.6 Ensure no unconfined services exist (Manual)

Profile Applicability:

- Level 1 - Server

Description:

Unconfined processes run in unconfined domains

Rationale:

For unconfined processes, SELinux policy rules are applied, but policy rules exist that allow processes running in unconfined domains almost all access. Processes running in unconfined domains fall back to using DAC rules exclusively. If an unconfined process is compromised, SELinux does not prevent an attacker from gaining access to system resources and data, but of course, DAC rules are still used. SELinux is a security enhancement on top of DAC rules – it does not replace them

Audit:

Run the following command and verify no output is produced:

```
# ps -eZ | grep unconfined_service_t
```

Nothing should be returned

Remediation:

Investigate any unconfined processes found during the audit action. They may need to have an existing security context assigned to them or a policy built for them.

References:

1. NIST SP 800-53 Rev. 5: AC-3, MP-2

Additional Information:

Occasionally certain daemons such as backup or centralized management software may require running unconfined. Any such software should be carefully analyzed and documented before such an exception is made.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	9.2 <u>Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1068, T1068.000, T1565, T1565.001, T1565.003	TA0004	M1022

1.6.1.7 Ensure SETroubleshoot is not installed (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The SETroubleshoot service notifies desktop users of SELinux denials through a user-friendly interface. The service provides important information around configuration errors, unauthorized intrusions, and other potential errors.

Rationale:

The SETroubleshoot service is an unnecessary daemon to have running on a server, especially if X Windows is disabled.

Audit:

Verify `setroubleshoot` is not installed.

Run the following command:

```
# rpm -q setroubleshoot  
package setroubleshoot is not installed
```

Remediation:

Run the following command to uninstall `setroubleshoot`:

```
# dnf remove setroubleshoot
```

References:

1. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 <u>Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1543, T1543.002	TA0005	M1033

1.6.1.8 Ensure the MCS Translation Service (mcstrans) is not installed (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `mcstransd` daemon provides category label information to client processes requesting information. The label translations are defined in `/etc/selinux/targeted/setrans.conf`

Rationale:

Since this service is not used very often, remove it to reduce the amount of potentially vulnerable code running on the system.

Audit:

Verify `mcstrans` is not installed.
Run the following command:

```
# rpm -q mcstrans  
package mcstrans is not installed
```

Remediation:

Run the following command to uninstall `mcstrans`:

```
# dnf remove mcstrans
```

References:

1. NIST SP 800-53 Rev. 5: SI-4

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1543, T1543.002	TA0005	M1033

1.7 Command Line Warning Banners

Presenting a warning message prior to the normal user login may assist in the prosecution of trespassers on the computer system. Changing some of these login banners also has the side effect of hiding OS version information and other detailed system information from attackers attempting to target specific exploits at a system.

Guidelines published by the US Department of Defense require that warning messages include at least the name of the organization that owns the system, the fact that the system is subject to monitoring and that such monitoring is in compliance with local statutes, and that use of the system implies consent to such monitoring. It is important that the organization's legal counsel review the content of all messages before any system modifications are made, as these warning messages are inherently site-specific. More information (including citations of relevant case law) can be found at <http://www.justice.gov/criminal/cybercrime/>

The `/etc/motd`, `/etc/issue`, and `/etc/issue.net` files govern warning banners for standard command line logins for both local and remote users.

Note: The text provided in the remediation actions for these items is intended as an example only. Please edit to include the specific text for your organization as approved by your legal department.

1.7.1 Ensure message of the day is configured properly (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The contents of the `/etc/motd` file are displayed to users after login and function as a message of the day for authenticated users.

Unix-based systems have typically displayed information about the OS release and patch level upon logging in to the system. This information can be useful to developers who are developing software for a particular OS platform. If `mingetty(8)` supports the following options, they display operating system information: `\m` - machine architecture `\r` - operating system release `\s` - operating system name `\v` - operating system version

Rationale:

Warning messages inform users who are attempting to login to the system of their legal status regarding the system and must include the name of the organization that owns the system and any monitoring policies that are in place. Displaying OS and patch level information in login banners also has the side effect of providing detailed system information to attackers attempting to target specific exploits of a system. Authorized users can easily get this information by running the "`uname -a`" command once they have logged in.

Audit:

Run the following command and verify that the contents match site policy:

```
# cat /etc/motd
```

Run the following command and verify no results are returned:

```
# grep -E -i "(\\v|\\r|\\m|\\s|$(grep '^ID=' /etc/os-release | cut -d= -f2 | sed -e 's/"/g'))" /etc/motd
```

Remediation:

Edit the `/etc/motd` file with the appropriate contents according to your site policy, remove any instances of `\m`, `\r`, `\s`, `\v` or references to the OS platform

OR

If the `motd` is not used, this file can be removed.

Run the following command to remove the `motd` file:

```
# rm /etc/motd
```

References:

1. NIST SP 800-53 Rev. 5: CM-3

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1082, T1082.000, T1592, T1592.004	TA0007	

1.7.2 Ensure local login warning banner is configured properly (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The contents of the `/etc/issue` file are displayed to users prior to login for local terminals.

Unix-based systems have typically displayed information about the OS release and patch level upon logging in to the system. This information can be useful to developers who are developing software for a particular OS platform. If `mingetty(8)` supports the following options, they display operating system information: `\m` - machine architecture `\r` - operating system release `\s` - operating system name `\v` - operating system version - or the operating system's name

Rationale:

Warning messages inform users who are attempting to login to the system of their legal status regarding the system and must include the name of the organization that owns the system and any monitoring policies that are in place. Displaying OS and patch level information in login banners also has the side effect of providing detailed system information to attackers attempting to target specific exploits of a system. Authorized users can easily get this information by running the "`uname -a`" command once they have logged in.

Audit:

Run the following command and verify that the contents match site policy:

```
# cat /etc/issue
```

Run the following command and verify no results are returned:

```
# grep -E -i "(\\v|\\r|\\m|\\s|$(grep '^ID=' /etc/os-release | cut -d= -f2 | sed -e 's/"/g'))" /etc/issue
```

Remediation:

Edit the `/etc/issue` file with the appropriate contents according to your site policy, remove any instances of `\m`, `\r`, `\s`, `\v` or references to the OS platform

```
# echo "Authorized uses only. All activity may be monitored and reported." > /etc/issue
```

References:

1. NIST SP 800-53 Rev. 5: CM-6, CM-1, CM-3

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1082, T1082.000, T1592, T1592.004	TA0007	

1.7.3 Ensure remote login warning banner is configured properly (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The contents of the `/etc/issue.net` file are displayed to users prior to login for remote connections from configured services.

Unix-based systems have typically displayed information about the OS release and patch level upon logging in to the system. This information can be useful to developers who are developing software for a particular OS platform. If `mingetty(8)` supports the following options, they display operating system information: `\m` - machine architecture `\r` - operating system release `\s` - operating system name `\v` - operating system version

Rationale:

Warning messages inform users who are attempting to login to the system of their legal status regarding the system and must include the name of the organization that owns the system and any monitoring policies that are in place. Displaying OS and patch level information in login banners also has the side effect of providing detailed system information to attackers attempting to target specific exploits of a system. Authorized users can easily get this information by running the "`uname -a`" command once they have logged in.

Audit:

Run the following command and verify that the contents match site policy:

```
# cat /etc/issue.net
```

Run the following command and verify no results are returned:

```
# grep -E -i "(\\v|\\r|\\m|\\s|$(grep '^ID=' /etc/os-release | cut -d= -f2 | sed -e 's/"/g'))" /etc/issue.net
```

Remediation:

Edit the `/etc/issue.net` file with the appropriate contents according to your site policy, remove any instances of `\m`, `\r`, `\s`, `\v` or references to the OS platform

```
# echo "Authorized uses only. All activity may be monitored and reported." > /etc/issue.net
```

References:

1. NIST SP 800-53 Rev. 5: CM-6, CM-1, CM-3

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1018, T1018.000, T1082, T1082.000, T1592, T1592.004	TA0007	

1.7.4 Ensure permissions on /etc/motd are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The contents of the `/etc/motd` file are displayed to users after login and function as a message of the day for authenticated users.

Rationale:

If the `/etc/motd` file does not have the correct ownership it could be modified by unauthorized users with incorrect or misleading information.

Audit:

Run the following command and verify `Uid` and `Gid` are both `0/root` and `Access` is `644` :

```
# stat -L /etc/motd
Access: (0644/-rw-r--r--)  Uid: (   0/   root)   Gid: (   0/   root)
```

Remediation:

Run the following commands to set permissions on `/etc/motd` :

```
# chown root:root /etc/motd
# chmod u-x,go-wx /etc/motd
```

References:

1. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1222, T1222.002	TA0005	M1022

1.7.5 Ensure permissions on /etc/issue are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The contents of the `/etc/issue` file are displayed to users prior to login for local terminals.

Rationale:

If the `/etc/issue` file does not have the correct ownership it could be modified by unauthorized users with incorrect or misleading information.

Audit:

Run the following command and verify `Uid` and `Gid` are both `0/root` and `Access` is `644` :

```
# stat -L /etc/issue
Access: (0644/-rw-r--r--)  Uid: (   0/   root)   Gid: (   0/   root)
```

Remediation:

Run the following commands to set permissions on `/etc/issue` :

```
# chown root:root /etc/issue
# chmod u-x,go-wx /etc/issue
```

References:

1. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1222, T1222.002	TA0005	M1022

1.7.6 Ensure permissions on /etc/issue.net are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The contents of the `/etc/issue.net` file are displayed to users prior to login for remote connections from configured services.

Rationale:

If the `/etc/issue.net` file does not have the correct ownership it could be modified by unauthorized users with incorrect or misleading information.

Audit:

Run the following command and verify `Uid` and `Gid` are both `0/root` and `Access` is `644` :

```
# stat -L /etc/issue.net
Access: (0644/-rw-r--r--)  Uid: (   0/   root)   Gid: (   0/   root)
```

Remediation:

Run the following commands to set permissions on `/etc/issue.net` :

```
# chown root:root /etc/issue.net
# chmod u-x,go-wx /etc/issue.net
```

References:

1. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1222, T1222.002	TA0005	M1022

1.8 Ensure updates, patches, and additional security software are installed (Manual)

Profile Applicability:

- Level 1 - Server

Description:

Periodically patches are released for included software either due to security flaws or to include additional functionality.

Rationale:

Newer patches may contain security enhancements that would not be available through the latest full update. As a result, it is recommended that the latest software patches be used to take advantage of the latest functionality. As with any software installation, organizations need to determine if a given update meets their requirements and verify the compatibility and supportability of any additional software against the update revision that is selected.

Audit:

Run the following command and verify there are no updates or patches to install:

```
# dnf check-update

WARNING:
  A newer release of "Amazon Linux" is available.

Available Versions:

Version 2023.0.20230503:
  Run the following command to upgrade to 2023.0.XXXXX:

    dnf upgrade --releasever=2023.0.XXXXX

Release notes:
  https://docs.aws.amazon.com/linux/al2023/release-notes/relnotes.html
```

Check to make sure no system reboot is required

```
dnf needs-restarting -r
```

Remediation:

Use your package manager to update all packages on the system according to site policy.

Based on the version available to install, run the following command with the correct versioning:

```
# dnf upgrade --releasever=2023.0.XXXXXXX
```

Once the update process is complete, verify if reboot is required to load changes.

```
dnf needs-restarting -r
```

References:

1. NIST SP 800-53 Rev. 5: SI-2

Additional Information:

Site policy may mandate a testing period before install onto production systems for available updates.

```
# dnf check-update
```

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>7.3 Perform Automated Operating System Patch Management</u> Perform operating system updates on enterprise assets through automated patch management on a monthly, or more frequent, basis.			
v8	<u>7.4 Perform Automated Application Patch Management</u> Perform application updates on enterprise assets through automated patch management on a monthly, or more frequent, basis.			
v7	<u>3.4 Deploy Automated Operating System Patch Management Tools</u> Deploy automated software update tools in order to ensure that the operating systems are running the most recent security updates provided by the software vendor.			
v7	<u>3.5 Deploy Automated Software Patch Management Tools</u> Deploy automated software update tools in order to ensure that third-party software on all systems is running the most recent security updates provided by the software vendor.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1211, T1211.000	TA0004, TA0008	M1051

1.9 Ensure system-wide crypto policy is not legacy (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The system-wide crypto-policies followed by the crypto core components allow consistently deprecating and disabling algorithms system-wide.

The individual policy levels (DEFAULT, LEGACY, FUTURE, and FIPS) are included in the crypto-policies(7) package.

Rationale:

If the Legacy system-wide crypto policy is selected, it includes support for TLS 1.0, TLS 1.1, and SSH2 protocols or later. The algorithms DSA, 3DES, and RC4 are allowed, while RSA and Diffie-Hellman parameters are accepted if larger than 1023-bits.

These legacy protocols and algorithms can make the system vulnerable to attacks, including those listed in RFC 7457

Impact:

Environments that require compatibility with older insecure protocols may require the use of the less secure LEGACY policy level.

Audit:

Run the following command to verify that the system-wide crypto policy is not LEGACY

```
# grep -E -i '^s*LEGACY\s*(\s+#.*)?$' /etc/crypto-policies/config
```

Verify that no lines are returned

Remediation:

Run the following command to change the system-wide crypto policy

```
# update-crypto-policies --set <CRYPTO POLICY>
```

Example:

```
# update-crypto-policies --set DEFAULT
```

Run the following to make the updated system-wide crypto policy active

```
# update-crypto-policies
```

Default Value:

DEFAULT

References:

1. CRYPTO-POLICIES(7)
2. <https://access.redhat.com/articles/3642912#what-polices-are-provided-1>
3. NIST SP 800-53 Rev. 5: SC-8

Additional Information:

To switch the system to FIPS mode, run the following command:

```
fips-mode-setup --enable
```

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.10 <u>Encrypt Sensitive Data in Transit</u> Encrypt sensitive data in transit. Example implementations can include: Transport Layer Security (TLS) and Open Secure Shell (OpenSSH).			
v7	14.4 <u>Encrypt All Sensitive Information in Transit</u> Encrypt all sensitive information in transit.			

2 Services

While applying system updates and patches helps correct known vulnerabilities, one of the best ways to protect the system against as yet unreported vulnerabilities is to disable all services that are not required for normal system operation. This prevents the exploitation of vulnerabilities discovered at a later date. If a service is not enabled, it cannot be exploited. The actions in this section of the document provide guidance on some services which can be safely disabled and under which circumstances, greatly reducing the number of possible threats to the resulting system. Additionally, some services which should remain enabled but with secure configuration are covered as well as insecure service clients.

2.1 Configure Time Synchronization

It is recommended that physical systems and virtual guests lacking direct access to the physical host's clock be configured to synchronize their time using a service such as NTP or chrony.

2.1.1 Ensure time synchronization is in use (Automated)

Profile Applicability:

- Level 1 - Server

Description:

System time should be synchronized between all systems in an environment. This is typically done by establishing an authoritative time server or set of servers and having all systems synchronize their clocks to them.

Note: If another method for time synchronization is being used, this section may be skipped.

Rationale:

Time synchronization is important to support time sensitive security mechanisms like Kerberos and also ensures log files have consistent time records across the enterprise, which aids in forensic investigations.

Audit:

Run the following commands to verify that chrony is installed:

```
# rpm -q chrony  
chrony-<version>
```

Remediation:

Run the following command to install `chrony`:

```
# dnf install chrony
```

References:

1. NIST SP 800-53 Rev. 5: AU-3, AU-12

Additional Information:

On systems where host based time synchronization is not available, verify that chrony is installed.

On systems where host based time synchronization is available consult your documentation and verify that host based synchronization is in use.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.4 <u>Standardize Time Synchronization</u> Standardize time synchronization. Configure at least two synchronized time sources across enterprise assets, where supported.		●	●
v7	6.1 <u>Utilize Three Synchronized Time Sources</u> Use at least three synchronized time sources from which all servers and network devices retrieve time information on a regular basis so that timestamps in logs are consistent.		●	●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1562, T1562.001	TA0005	M1022

2.1.2 Ensure chrony is configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

`chrony` is a daemon which implements the Network Time Protocol (NTP) and is designed to synchronize system clocks across a variety of systems and use a source that is highly accurate. More information on `chrony` can be found at <http://chrony.tuxfamily.org/>. `chrony` can be configured to be a client and/or a server.

Rationale:

If `chrony` is in use on the system proper configuration is vital to ensuring time synchronization is working properly.

Audit:

Run the following command and verify remote server is configured properly:

```
# grep -Pr -- '^h*(server|pool)\h+[^#\n\r]+' /etc/chrony.d/  
server <remote-server>
```

Multiple servers may be configured.

Remediation:

Add or edit server or pool lines to file ending in `.conf` in the `/etc/chrony.d` as appropriate:

Example:

```
server <remote-server>
```

References:

1. NIST SP 800-53 Rev. 5: AU-3, AU-12

Additional Information:

On systems where host based time synchronization is not available, verify that `chrony` is installed.

On systems where host based time synchronization is available consult your documentation and verify that host based synchronization is in use.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.4 <u>Standardize Time Synchronization</u> Standardize time synchronization. Configure at least two synchronized time sources across enterprise assets, where supported.		●	●
v7	6.1 <u>Utilize Three Synchronized Time Sources</u> Use at least three synchronized time sources from which all servers and network devices retrieve time information on a regular basis so that timestamps in logs are consistent.		●	●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002	TA0002	M1022

2.2 Configure Special Purpose Services

This section describes services that are installed on systems that specifically need to run these services. If any of these services are not required, it is recommended that the package be removed, or the service be masked to reduce the potential attack surface.

Note: This should not be considered a comprehensive list of services not required for normal system operation. You may wish to consider additions to those listed here for your environment

2.2.1 Ensure xorg-x11-server-common is not installed (Automated)

Profile Applicability:

- Level 2 - Server

Description:

The X Window System provides a Graphical User Interface (GUI) where users can have multiple windows in which to run programs and various add on. The X Windows system is typically used on workstations where users login, but not on servers where users typically do not login.

Rationale:

Unless your organization specifically requires graphical login access via X Windows, remove it to reduce the potential attack surface.

Impact:

Many Linux systems run applications which require a Java runtime. Some Linux Java packages have a dependency on specific X Windows xorg-x11-fonts. One workaround to avoid this dependency is to use the "headless" Java packages for your specific Java runtime.

Audit:

Run the following command to Verify X Windows Server is not installed.

```
# rpm -q xorg-x11-server-common  
package xorg-x11-server-common is not installed
```

Remediation:

Run the following command to remove the X Windows Server packages:

```
# dnf remove xorg-x11-server-common
```

References:

1. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1203, T1203.000, T1210, T1210.000, T1543, T1543.002	TA0008	M1042

2.2.2 Ensure avahi is not installed (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Avahi is a free zeroconf implementation, including a system for multicast DNS/DNS-SD service discovery. Avahi allows programs to publish and discover services and hosts running on a local network with no specific configuration. For example, a user can plug a computer into a network and Avahi automatically finds printers to print to, files to look at and people to talk to, as well as network services running on the machine.

Rationale:

Automatic discovery of network services is not normally required for system functionality. It is recommended to remove this package to reduce the potential attack surface.

Audit:

Run one of the following command to verify `avahi` is not installed:

```
# rpm -q avahi  
  
package avahi is not installed
```

Remediation:

Run the following commands to stop, and remove `avahi`:

```
# systemctl stop avahi-daemon.socket avahi-daemon.service  
# dnf remove avahi
```

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●
v7	9.2 Ensure Only Approved Ports, Protocols and Services Are Running Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.		●	●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1203, T1203.000, T1210, T1210.000, T1543, T1543.002	TA0008	M1042

2.2.3 Ensure a print server is not installed (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The Common Unix Print System (CUPS) provides the ability to print to both local and network printers. A system running CUPS can also accept print jobs from remote systems and print them to local printers. It also provides a web based remote administration capability.

Rationale:

If the system does not need to print jobs or accept print jobs from other systems, it is recommended that CUPS be removed to reduce the potential attack surface.

Note: Removing CUPS will prevent printing from the system

Impact:

Disabling CUPS will prevent printing from the system, a common task for workstation systems.

Audit:

Run the following command to verify `cups` is not installed:

```
# rpm -q cups  
package cups is not installed
```

Remediation:

Run the following command to remove `cups`:

```
# dnf remove cups
```

References:

1. More detailed documentation on CUPS is available at the project homepage at <http://www.cups.org>.
2. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1203, T1203.000, T1210, T1210.000, T1543, T1543.002	TA0008	M1042

2.2.4 Ensure a dhcp server is not installed (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The Dynamic Host Configuration Protocol (DHCP) is a service that allows machines to be dynamically assigned IP addresses.

Rationale:

Unless a system is specifically set up to act as a DHCP server, it is recommended that the `dhcp-server` package be removed to reduce the potential attack surface.

Audit:

Run the following command to verify `dhcp-server` is not installed:

```
# rpm -q dhcp-server  
  
package dhcp-server is not installed
```

Remediation:

Run the following command to remove `dhcp`:

```
# dnf remove dhcp-server
```

References:

1. `dhcpcd(8)`
2. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 <u>Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	9.2 <u>Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1203, T1203.000, T1210, T1210.000, T1543, T1543.002	TA0008	M1042

2.2.5 Ensure a dns server is not installed (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The Domain Name System (DNS) is a hierarchical naming system that maps names to IP addresses for computers, services and other resources connected to a network.

Rationale:

Unless a system is specifically designated to act as a DNS server, it is recommended that the package be removed to reduce the potential attack surface.

Audit:

Run one of the following commands to verify `bind` is not installed:

```
# rpm -q bind  
  
package bind is not installed
```

Remediation:

Run the following command to remove `bind`:

```
# dnf remove bind
```

References:

1. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 <u>Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●
v7	9.2 <u>Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.		●	●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1203, T1203.000, T1210, T1210.000, T1543, T1543.002	TA0008	M1042

2.2.6 Ensure an ftp server is not installed (Automated)

Profile Applicability:

- Level 1 - Server

Description:

FTP (File Transfer Protocol) is a traditional and widely used standard tool for transferring files between a server and clients over a network, especially where no authentication is necessary (permits anonymous users to connect to a server).

Rationale:

Unless there is a need to run the system as a FTP server, it is recommended that the package be removed to reduce the potential attack surface.

Audit:

Run the following command to verify `vsftpd` is not installed:

```
# rpm -q vsftpd  
  
package vsftpd is not installed
```

Remediation:

Run the following command to remove `vsftpd`:

```
# dnf remove vsftpd
```

References:

1. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 <u>Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	9.2 <u>Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1203.000, T1210, T1210.000, T1543, T1543.002	TA0008	M1042

2.2.7 Ensure a tftp server is not installed (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Trivial File Transfer Protocol (TFTP) is a simple protocol for exchanging files between two TCP/IP machines. TFTP servers allow connections from a TFTP Client for sending and receiving files.

Rationale:

Unless there is a need to run the system as a TFTP server, it is recommended that the package be removed to reduce the potential attack surface.

TFTP does not have built-in encryption, access control or authentication. This makes it very easy for an attacker to exploit TFTP to gain access to files

Impact:

TFTP is often used to provide files for network booting such as for PXE based installation of servers.

Audit:

Run the following command to verify `tftp-server` is not installed:

```
# rpm -q tftp-server  
package tftp-server is not installed
```

Remediation:

Run the following command to remove `tftp-server`:

```
# dnf remove tftp-server
```

References:

1. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1203, T1203.000, T1210, T1210.000, T1543, T1543.002	TA0008	M1042

2.2.8 Ensure a web server is not installed (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Web servers provide the ability to host web site content.

Rationale:

Unless there is a need to run the system as a web server, it is recommended that the packages be removed to reduce the potential attack surface.

Note: Several http servers exist. They should also be audited, and removed, if not required.

Audit:

Run the following command to verify `httpd` and `nginx` are not installed:

```
# rpm -q httpd nginx  
  
package httpd is not installed  
package nginx is not installed
```

Remediation:

Run the following command to remove `httpd` and `nginx`:

```
# dnf remove httpd nginx
```

References:

1. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1203, T1203.000, T1210, T1210.000, T1543, T1543.002	TA0008	M1042

2.2.9 Ensure IMAP and POP3 server is not installed (Automated)

Profile Applicability:

- Level 1 - Server

Description:

`dovecot` is an open source IMAP and POP3 server for Linux based systems.

Rationale:

Unless POP3 and/or IMAP servers are to be provided by this system, it is recommended that the package be removed to reduce the potential attack surface.

Note: Several IMAP/POP3 servers exist and can use other service names. These should also be audited and the packages removed if not required.

Audit:

Run the following command to verify `dovecot` and `cyrus-imapd` are not installed:

```
# rpm -q dovecot cyrus-imapd  
  
package dovecot is not installed  
package cyrus-imapd is not installed
```

Remediation:

Run the following command to remove `dovecot` and `cyrus-imapd`:

```
# dnf remove dovecot cyrus-imapd
```

References:

1. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1203, T1203.000, T1210, T1210.000, T1543, T1543.002	TA0008	M1042

2.2.10 Ensure Samba is not installed (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The Samba daemon allows system administrators to configure their Linux systems to share file systems and directories with Windows desktops. Samba will advertise the file systems and directories via the Server Message Block (SMB) protocol. Windows desktop users will be able to mount these directories and file systems as letter drives on their systems.

Rationale:

If there is no need to mount directories and file systems to Windows systems, then this package can be removed to reduce the potential attack surface.

Audit:

Run the following command to verify `samba` is not installed:

```
# rpm -q samba  
package samba is not installed
```

Remediation:

Run the following command to remove `samba`:

```
# dnf remove samba
```

References:

1. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1005, T1005.000, T1039, T1039.000, T1083, T1083.000, T1135, T1135.000, T1203, T1203.000, T1210, T1210.000, T1543, T1543.002	TA0008	M1042

2.2.11 Ensure HTTP Proxy Server is not installed (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Squid is a standard proxy server used in many distributions and environments.

Rationale:

Unless a system is specifically set up to act as a proxy server, it is recommended that the squid package be removed to reduce the potential attack surface.

Note: Several HTTP proxy servers exist. These should be checked and removed unless required.

Audit:

Run the following command to verify `squid` is not installed:

```
# rpm -q squid  
  
package squid is not installed
```

Remediation:

Run the following command to remove the `squid` package:

```
# dnf remove squid
```

References:

1. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 <u>Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	9.2 <u>Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1203, T1203.000, T1210, T1210.000, T1543, T1543.002	TA0008	M1042

2.2.12 Ensure net-snmp is not installed or the snmpd service is not enabled (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Simple Network Management Protocol (SNMP) is a widely used protocol for monitoring the health and welfare of network equipment, computer equipment and devices like UPSs.

Net-SNMP is a suite of applications used to implement SNMPv1 (RFC 1157), SNMPv2 (RFCs 1901-1908), and SNMPv3 (RFCs 3411-3418) using both IPv4 and IPv6.

Support for SNMPv2 classic (a.k.a. "SNMPv2 historic" - RFCs 1441-1452) was dropped with the 4.0 release of the UCD-snmp package.

The Simple Network Management Protocol (SNMP) server is used to listen for SNMP commands from an SNMP management system, execute the commands or collect the information and then send results back to the requesting system.

Rationale:

The SNMP server can communicate using `SNMPv1`, which transmits data in the clear and does not require authentication to execute commands. `SNMPv3` replaces the simple/clear text password sharing used in `SNMPv2` with more securely encoded parameters. If the the SNMP service is not required, the `net-snmp` package should be removed to reduce the attack surface of the system.

Note: If a required dependency exists for the net-snmp package, but the snmpd service is not required, the service should be masked.

Note: If SNMP is required:

- The server should be configured for `SNMP v3` only. User Authentication and Message Encryption should be configured.
- If `SNMP v2` is **absolutely** necessary, modify the community strings' values.

Impact:

There are packages that are dependent on the net-snmp package. If the net-snmp package is removed, these packages will be removed as well.

Before removing the net-snmp package, review any dependent packages to determine if they are required on the system. If a dependent package is required, mask the snmpd service and leave the net-snmp package installed.

Audit:

Run the following command to verify `net-snmp` is not installed:

```
# rpm -q net-snmp  
  
package net-snmp is not installed
```

-OR-

Run the following command to verify the `snmpd` service is not enabled:

```
# systemctl is-enabled snmpd  
  
masked
```

Verify output is not `enabled`

Remediation:

Run the following command to remove `net-snmpd`:

```
# dnf remove net-snmp
```

-OR-

Run the following commands to stop and mask the `snmpd` service:

```
# systemctl stop snmpd  
# systemctl mask snmpd
```

References:

1. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 <u>Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	9.2 <u>Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1203, T1203.000, T1210, T1210.000, T1543, T1543.002	TA0008	M1042

2.2.13 Ensure telnet-server is not installed (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `telnet-server` package contains the `telnet` daemon, which accepts connections from users from other systems via the `telnet` protocol.

Rationale:

The `telnet` protocol is insecure and unencrypted. The use of an unencrypted transmission medium could allow a user with access to sniff network traffic the ability to steal credentials. The `ssh` package provides an encrypted session and stronger security.

Audit:

Run the following command to verify the `telnet-server` package is not installed:

```
rpm -q telnet-server  
package telnet-server is not installed
```

Remediation:

Run the following command to remove the `telnet-server` package:

```
# dnf remove telnet-server
```

References:

1. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●
v7	<u>2.6 Address unapproved software</u> Ensure that unauthorized software is either removed or the inventory is updated in a timely manner	●	●	●
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.		●	●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1203, T1203.000, T1210, T1210.000, T1543, T1543.002	TA0008	M1042

2.2.14 Ensure dnsmasq is not installed (Automated)

Profile Applicability:

- Level 1 - Server

Description:

`dnsmasq` is a lightweight tool that provides DNS caching, DNS forwarding and DHCP (Dynamic Host Configuration Protocol) services.

Rationale:

Unless a system is specifically designated to act as a DNS caching, DNS forwarding and/or DHCP server, it is recommended that the package be removed to reduce the potential attack surface.

Audit:

Run one of the following commands to verify `dnsmasq` is not installed:

```
# rpm -q dnsmasq  
package dnsmasq is not installed
```

Remediation:

Run the following command to remove `dnsmasq`:

```
# dnf remove dnsmasq
```

References:

1. NIST SP 800-53 Rev. 5: CM-7

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1203, T1203.000, T1210, T1210.000, T1543, T1543.002	TA0008	M1042

2.2.15 Ensure mail transfer agent is configured for local-only mode (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Mail Transfer Agents (MTA), such as sendmail and Postfix, are used to listen for incoming mail and transfer the messages to the appropriate user or mail server. If the system is not intended to be a mail server, it is recommended that the MTA be configured to only process local mail.

Rationale:

The software for all Mail Transfer Agents is complex and most have a long history of security issues. While it is important to ensure that the system can process local mail messages, it is not necessary to have the MTA's daemon listening on a port unless the server is intended to be a mail server that receives and processes mail from other systems.

Note:

- This recommendation is designed around the postfix mail server.
- Depending on your environment you may have an alternative MTA installed such as sendmail. If this is the case consult the documentation for your installed MTA to configure the recommended state.

Audit:

Run the following command to verify that the MTA is not listening on any non-loopback address (127.0.0.1 or ::1)
Nothing should be returned

```
# ss -ltnu | grep -P ':25\b' | grep -Pv '\h+(127\.0\.0\.1|\\[?::1\\]):25\b'
```

Remediation:

Edit `/etc/postfix/main.cf` and add the following line to the RECEIVING MAIL section. If the line already exists, change it to look like the line below:

```
inet_interfaces = loopback-only
```

Run the following command to restart postfix:

```
# systemctl restart postfix
```

References:

1. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	9.2 Ensure Only Approved Ports, Protocols and Services Are Running Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1018, T1018.000, T1210, T1210.000	TA0008	M1042

2.2.16 Ensure nfs-utils is not installed or the nfs-server service is masked (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The Network File System (NFS) is one of the first and most widely distributed file systems in the UNIX environment. It provides the ability for systems to mount file systems of other servers through the network.

Rationale:

If the system does not require network shares, it is recommended that the `nfs-utils` package be removed to reduce the attack surface of the system.

Impact:

Many of the `libvirt` packages used by Enterprise Linux virtualization are dependent on the `nfs-utils` package. If the `nfs-utils` package is required as a dependency, the `nfs-server` service should be disabled and masked to reduce the attack surface of the system.

Audit:

Run the following command to verify `nfs-utils` is not installed:

```
# rpm -q nfs-utils
package nfs-utils is not installed
```

OR

If the `nfs-utils` package is required as a dependency, run the following command to verify that the `nfs-server` service is masked:

```
# systemctl is-enabled nfs-server
masked
```

Remediation:

Run the following command to remove `nfs-utils`:

```
# dnf remove nfs-utils
```

OR

If the `nfs-utils` package is required as a dependency, run the following commands to stop and mask the `nfs-server` service:

```
# systemctl stop nfs-server  
# systemctl mask nfs-server
```

References:

1. NIST SP 800-53 Rev. 5: CM-6, CM-7

Additional Information:

Many of the libvirt packages used by Enterprise Linux virtualization are dependent on the `nfs-utils` package. If the `nfs-utils` package is required as a dependency, the `nfs-server` service should be disabled and masked to reduce the attack surface of the system.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 <u>Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●
v7	9.2 <u>Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.		●	●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1005, T1005.000, T1039, T1039.000, T1083, T1083.000, T1135, T1135.000, T1210, T1210.000	TA0008	M1042

2.2.17 Ensure rpcbind is not installed or the rpcbind services are masked (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The rpcbind utility maps RPC services to the ports on which they listen. RPC processes notify rpcbind when they start, registering the ports they are listening on and the RPC program numbers they expect to serve. The client system then contacts rpcbind on the server with a particular RPC program number. The rpcbind service redirects the client to the proper port number so it can communicate with the requested service.

Portmapper is an RPC service, which always listens on tcp and udp 111, and is used to map other RPC services (such as nfs, nlockmgr, quotad, mountd, etc.) to their corresponding port number on the server. When a remote host makes an RPC call to that server, it first consults with portmap to determine where the RPC server is listening.

Rationale:

A small request (~82 bytes via UDP) sent to the Portmapper generates a large response (7x to 28x amplification), which makes it a suitable tool for DDoS attacks. If rpcbind is not required, it is recommended that the rpcbind package be removed to reduce the attack surface of the system.

Impact:

Many of the libvirt packages used by Enterprise Linux virtualization, and the `nfs-utils` package used for The Network File System (NFS), are dependent on the `rpcbind` package. If the `rpcbind` package is required as a dependency, the services `rpcbind.service` and `rpcbind.socket` should be stopped and masked to reduce the attack surface of the system.

Audit:

Run the following command to verify `rpcbind` is not installed:

```
# rpm -q rpcbind  
  
package rpcbind is not installed
```

OR

If the `rpcbind` package is required as a dependency, run the following commands to verify that the `rpcbind` and `rpcbind.socket` `systemd` units are masked:

```
# systemctl is-enabled rpcbind  
  
masked  
# systemctl is-enabled rpcbind.socket  
  
masked
```

Remediation:

Run the following command to remove `nfs-utils`:

```
# dnf remove rpcbind
```

OR

If the `rpcbind` package is required as a dependency, run the following commands to stop and mask the `rpcbind.service` and `rpcbind.socket` `systemd` units:

```
# systemctl stop rpcbind.socket  
# systemctl mask rpcbind.socket  
  
# systemctl stop rpcbind.service  
# systemctl mask rpcbind.service
```

References:

1. NIST SP 800-53 Rev. 5: CM-6, CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1203, T1203.000, T1210, T1210.000, T1498, T1498.002, T1543, T1543.002	TA0008	M1042

2.2.18 Ensure rsync-daemon is not installed or the rsyncd service is masked (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `rsyncd` service can be used to synchronize files between systems over network links.

Rationale:

Unless required, the `rsync-daemon` package should be removed to reduce the attack surface area of the system.

The `rsyncd` service presents a security risk as it uses unencrypted protocols for communication.

Note: If a required dependency exists for the `rsync-daemon` package, but the `rsyncd` service is not required, the service should be masked.

Impact:

There are packages that are dependent on the `rsync` package. If the `rsync` package is removed, these packages will be removed as well.

Before removing the `rsync-daemon` package, review any dependent packages to determine if they are required on the system. If a dependent package is required, mask the `rsyncd` service and leave the `rsync-daemon` package installed.

Audit:

Run the following command to verify that `rsync` is not installed:

```
# rpm -q rsync-daemon
package rsync is not installed
```

OR

Run the following commands to verify the `rsyncd` service is masked:

```
# systemctl is-enabled rsyncd
masked
```

Remediation:

Run the following command to remove the `rsync` package:

```
# dnf remove rsync-daemon
```

OR

Run the following commands to mask the `rsyncd` service:

```
# systemctl stop rsyncd  
# systemctl mask rsyncd
```

References:

1. NIST SP 800-53 Rev. 5: CM-6, CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1105, T1105.000, T1203, T1203.000, T1210, T1210.000, T1543, T1543.002, T1570, T1570.000	TA0008	M1042

2.3 Service Clients

A number of insecure services exist. While disabling the servers prevents a local attack against these services, it is advised to remove their clients unless they are required.

Note: This should not be considered a comprehensive list of insecure service clients. You may wish to consider additions to those listed here for your environment.

2.3.1 Ensure telnet client is not installed (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `telnet` package contains the `telnet` client, which allows users to start connections to other systems via the telnet protocol.

Rationale:

The `telnet` protocol is insecure and unencrypted. The use of an unencrypted transmission medium could allow an unauthorized user to steal credentials. The `ssh` package provides an encrypted session and stronger security and is included in most Linux distributions.

Impact:

Many insecure service clients are used as troubleshooting tools and in testing environments. Uninstalling them can inhibit capability to test and troubleshoot. If they are required it is advisable to remove the clients after use to prevent accidental or intentional misuse.

Audit:

Run the following command to verify that the `telnet` package is not installed:

```
# rpm -q telnet  
package telnet is not installed
```

Remediation:

Run the following command to remove the `telnet` package:

```
# dnf remove telnet
```

References:

1. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 <u>Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	2.6 <u>Address unapproved software</u> Ensure that unauthorized software is either removed or the inventory is updated in a timely manner			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1040, T1040.000, T1203, T1203.000, T1543, T1543.002	TA0006, TA0008	M1041, M1042

2.3.2 Ensure LDAP client is not installed (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The Lightweight Directory Access Protocol (LDAP) was introduced as a replacement for NIS/YP. It is a service that provides a method for looking up information from a central database.

Rationale:

If the system will not need to act as an LDAP client, it is recommended that the software be removed to reduce the potential attack surface.

Impact:

Removing the LDAP client will prevent or inhibit using LDAP for authentication in your environment.

Audit:

Run the following command to verify that the `openldap-clients` package is not installed:

```
# rpm -q openldap-clients  
package openldap-clients is not installed
```

Remediation:

Run the following command to remove the `openldap-clients` package:

```
# dnf remove openldap-clients
```

References:

1. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>2.6 Address unapproved software</u> Ensure that unauthorized software is either removed or the inventory is updated in a timely manner			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1203, T1203.000, T1543, T1543.002	TA0008	M1042

2.3.3 Ensure FTP client is not installed (Automated)

Profile Applicability:

- Level 1 - Server

Description:

FTP (File Transfer Protocol) is a traditional and widely used standard tool for transferring files between a server and clients over a network, especially where no authentication is necessary (permits anonymous users to connect to a server).

Rationale:

FTP does not protect the confidentiality of data or authentication credentials. It is recommended SFTP be used if file transfer is required. Unless there is a need to run the system as a FTP server (for example, to allow anonymous downloads), it is recommended that the package be removed to reduce the potential attack surface.

Audit:

Run the following command to verify `ftp` is not installed:

```
# rpm -q ftp  
package ftp is not installed
```

Remediation:

Run the following command to remove `ftp`:

```
# dnf remove ftp
```

References:

1. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1203, T1203.000, T1543, T1543.002	TA0008	M1042

2.4 Ensure nonessential services listening on the system are removed or masked (Manual)

Profile Applicability:

- Level 1 - Server

Description:

A network port is identified by its number, the associated IP address, and the type of the communication protocol such as TCP or UDP.

A listening port is a network port on which an application or process listens on, acting as a communication endpoint.

Each listening port can be open or closed (filtered) using a firewall. In general terms, an open port is a network port that accepts incoming packets from remote locations.

Rationale:

Services listening on the system pose a potential risk as an attack vector. These services should be reviewed, and if not required, the service should be stopped, and the package containing the service should be removed. If required packages have a dependency, the service should be stopped and masked to reduce the attack surface of the system.

Audit:

Run the following command:

```
# ss -plntu
```

Review the output to ensure that all services listed are required on the system. If a listed service is not required, remove the package containing the service. If the package containing the service is required, stop and mask the service

Remediation:

Run the following command to remove the package containing the service:

```
# dnf remove <package_name>
```

OR If required packages have a dependency:

Run the following commands to stop and mask the service:

```
# systemctl stop <service_name>.socket
# systemctl stop <service_name>.service
# systemctl mask <service_name>.socket
# systemctl mask <service_name>.service
```

References:

1. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 <u>Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	9.2 <u>Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1203, T1203.000, T1210, T1210.000, T1543, T1543.002	TA0008	M1042

3 Network Configuration

This section provides guidance on for securing the network configuration of the system through kernel parameters, access list control, and firewall settings.

Note:

- sysctl settings are defined through files in `/usr/lib/sysctl.d/`, `/run/sysctl.d/`, and `/etc/sysctl.d/`.
- Files must have the `".conf"` extension.
- Vendors settings live in `/usr/lib/sysctl.d/`
- To override a whole file, create a new file with the same name in `/etc/sysctl.d/` and put new settings there.
- To override only specific settings, add a file with a lexically later name in `/etc/sysctl.d/` and put new settings there.
- The paths where sysctl preload files usually exist
 - `/run/sysctl.d/*.conf`
 - `/etc/sysctl.d/*.conf`
 - `/usr/local/lib/sysctl.d/*.conf`
 - `/usr/lib/sysctl.d/*.conf`
 - `/lib/sysctl.d/*.conf`
 - `/etc/sysctl.conf`

3.1 Disable unused network protocols and devices

To reduce the attack surface of a system, unused network protocols and devices should be disabled.

The Linux kernel modules support several network protocols that are not commonly used. If these protocols are not needed, it is recommended that they be disabled in the kernel.

3.1.1 Ensure IPv6 status is identified (Manual)

Profile Applicability:

- Level 1 - Server

Description:

Internet Protocol Version 6 (IPv6) is the most recent version of Internet Protocol (IP). It's designed to supply IP addressing and additional security to support the predicted growth of connected devices. IPv6 is based on 128-bit addressing and can support 340 undecillion addresses, which is 340 followed by 36 zeroes.

Features of IPv6

- Hierarchical addressing and routing infrastructure
- Stateful and Stateless configuration
- Support for quality of service (QoS)
- An ideal protocol for neighboring node interaction

Rationale:

IETF RFC 4038 recommends that applications are built with an assumption of dual stack. It is recommended that IPv6 be enabled and configured in accordance with Benchmark recommendations.

If dual stack and IPv6 are not used in your environment, IPv6 may be disabled to reduce the attack surface of the system, and recommendations pertaining to IPv6 can be skipped.

Note: It is recommended that IPv6 be enabled and configured unless this is against local site policy

Impact:

IETF RFC 4038 recommends that applications are built with an assumption of dual stack.

When enabled, IPv6 will require additional configuration to reduce risk to the system.

Audit:

Run the following to identify if IPv6 is enabled on the system:

```
# grep -Pqs '^h*0\b' /sys/module/ipv6/parameters/disable && echo -e "\n - IPv6 is enabled\n" || echo -e "\n - IPv6 is not enabled\n"
```

Remediation:

Enable or disable IPv6 in accordance with system requirements and local site policy

Default Value:

IPv6 is enabled

References:

1. NIST SP 800-53 Rev. 5: CM-7

Additional Information:

Having more addresses has grown in importance with the expansion of smart devices and connectivity. IPv6 provides more than enough globally unique IP addresses for every networked device currently on the planet, helping ensure providers can keep pace with the expected proliferation of IP-based devices.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.		●	●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1557, T1557.000, T1595, T1595.001, T1595.002	TA0008	M1042

3.1.2 Ensure DCCP is disabled (Automated)

Profile Applicability:

- Level 2 - Server

Description:

The Datagram Congestion Control Protocol (DCCP) is a transport layer protocol that supports streaming media and telephony. DCCP provides a way to gain access to congestion control, without having to do it at the application layer, but does not provide in-sequence delivery.

Rationale:

If the protocol is not required, it is recommended that the drivers not be installed to reduce the potential attack surface.

Audit:

Run the following script to verify the `dccp` module is disabled:

If the module is available in the running kernel:

- An entry including `/bin/true` or `/bin/false` exists in a file within the `/etc/modprobe.d/` directory
- The module is deny listed in a file within the `/etc/modprobe.d/` directory
- The module is not loaded in the kernel

If available in ANY installed kernel:

- The module is deny listed in a file within the `/etc/modprobe.d/` directory

If the kernel module is not available on the system, or pre-compiled into the kernel:

- No additional configuration is necessary


```

#!/usr/bin/env bash

{
    l_output="" l_output2="" l_output3="" l_dl="" # Unset output variables
    l_mname="dccc" # set module name
    l_mtype="net" # set module type
    l_searchloc="/lib/modprobe.d/*.conf /usr/local/lib/modprobe.d/*.conf /run/modprobe.d/*.conf
/etc/modprobe.d/*.conf"
    l_mpath="/lib/modules/**/kernel/$l_mtype"
    l_mpname="$(tr '-' '/' <<< "$l_mname")"
    l_mndir="$(tr '-' '/' <<< "$l_mname")"

    module_loadable_chk()
    {
        # Check if the module is currently loadable
        l_loadable="$(modprobe -n -v "$l_mname")"
        [ "$l_loadable" -gt "1" ] && l_loadable="$(grep -P --
"^(^\\h*install|\\b$l_mname)\\b" <<< "$l_loadable")"
        if grep -Pq -- '^\\h*install \\bin/(true|false)' <<< "$l_loadable"; then
            l_output="$l_output\\n - module: \"$l_mname\" is not loadable: \"$l_loadable\\n\"
        else
            l_output2="$l_output2\\n - module: \"$l_mname\" is loadable: \"$l_loadable\\n\"
        fi
    }

    module_loaded_chk()
    {
        # Check if the module is currently loaded
        if ! lsmod | grep "$l_mname" > /dev/null 2>&1; then
            l_output="$l_output\\n - module: \"$l_mname\" is not loaded"
        else
            l_output2="$l_output2\\n - module: \"$l_mname\" is loaded"
        fi
    }

    module_deny_chk()
    {
        # Check if the module is deny listed
        l_dl="y"
        if modprobe --showconfig | grep -Pq -- '^\\h*blacklist\\h+'$l_mpname'"\\b'; then
            l_output="$l_output\\n - module: \"$l_mname\" is deny listed in: \"$(grep -Pls --
"^(^\\h*blacklist\\h+$l_mname\\b" $l_searchloc)\"
        else
            l_output2="$l_output2\\n - module: \"$l_mname\" is not deny listed"
        fi
    }

    # Check if the module exists on the system
    for l_mdir in $l_mpath; do
        if [ -d "$l_mdir/$l_mndir" ] && [ -n "$(ls -A $l_mdir/$l_mndir)" ]; then
            l_output3="$l_output3\\n - \"$l_mdir\\n\"
            [ "$l_dl" != "y" ] && module_deny_chk
            if [ "$l_mdir" = "/lib/modules/$(uname -r)/kernel/$l_mtype" ]; then
                module_loadable_chk
                module_loaded_chk
            fi
        else
            l_output="$l_output\\n - module: \"$l_mname\" doesn't exist in \"$l_mdir\\n\"
        fi
    done

    # Report results. If no failures output in l_output2, we pass
    [ -n "$l_output3" ] && echo -e "\\n\\n -- INFO --\\n - module: \"$l_mname\" exists in:$l_output3"
    if [ -z "$l_output2" ]; then
        echo -e "\\n- Audit Result:\\n ** PASS **\\n$l_output\\n"
    else
        echo -e "\\n- Audit Result:\\n ** FAIL **\\n - Reason(s) for audit failure:\\n$l_output2\\n"
        [ -n "$l_output" ] && echo -e "\\n- Correctly set:\\n$l_output\\n"
    fi
}

```

Remediation:

Run the following script to disable the `dccp` module:

If the module is available in the running kernel:

- Create a file with `install dccp /bin/false` in the `/etc/modprobe.d/` directory
- Create a file with `blacklist dccp` in the `/etc/modprobe.d/` directory
- Unload `dccp` from the kernel

If available in ANY installed kernel:

- Create a file with `blacklist dccp` in the `/etc/modprobe.d/` directory

If the kernel module is not available on the system or pre-compiled into the kernel:

- No remediation is necessary

```
#!/usr/bin/env bash

{
  l_mname="dccc" # set module name
  l_mtype="net" # set module type
  l_mpath="/lib/modules/**/kernel/$l_mtype"
  l_mpname="$(tr '-' '_' <<< "$l_mname")"
  l_mndir="$(tr '-' '/' <<< "$l_mname")"

  module_loadable_fix()
  {
    # If the module is currently loadable, add "install {MODULE_NAME} /bin/false" to a file in
    "/etc/modprobe.d"
    l_loadable="$(modprobe -n -v "$l_mname")"
    [ "$(wc -l <<< "$l_loadable")" -gt "1" ] && l_loadable="$(grep -P --
    "(^\\h*install|\\b$l_mname)\\b" <<< "$l_loadable")"
    if ! grep -Pq -- '^\\h*install \\bin/(true|false)' <<< "$l_loadable"; then
      echo -e "\\n - setting module: \"$l_mname\" to be not loadable"
      echo -e "install $l_mname /bin/false" >> /etc/modprobe.d/"$l_mpname".conf
    fi
  }
  module_loaded_fix()
  {
    # If the module is currently loaded, unload the module
    if lsmod | grep "$l_mname" > /dev/null 2>&1; then
      echo -e "\\n - unloading module \"$l_mname\""
      modprobe -r "$l_mname"
    fi
  }
  module_deny_fix()
  {
    # If the module isn't deny listed, denylist the module
    if ! modprobe --showconfig | grep -Pq -- "^\\h*blacklist\\h+$l_mpname\\b"; then
      echo -e "\\n - deny listing \"$l_mname\""
      echo -e "blacklist $l_mname" >> /etc/modprobe.d/"$l_mpname".conf
    fi
  }
  # Check if the module exists on the system
  for l_mdir in $l_mpath; do
    if [ -d "$l_mdir/$l_mndir" ] && [ -n "$(ls -A $l_mdir/$l_mndir)" ]; then
      echo -e "\\n - module: \"$l_mname\" exists in \"$l_mdir\"\\n - checking if disabled..."
      module_deny_fix
      if [ "$l_mdir" = "/lib/modules/$(uname -r)/kernel/$l_mtype" ]; then
        module_loadable_fix
        module_loaded_fix
      fi
    else
      echo -e "\\n - module: \"$l_mname\" doesn't exist in \"$l_mdir\"\\n"
    fi
  done
  echo -e "\\n - remediation of module: \"$l_mname\" complete\\n"
}

```

References:

1. NIST SP 800-53 Rev. 5: SI-4, CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1068, T1068.000, T1210, T1210.000	TA0008	M1042

3.1.3 Ensure SCTP is disabled (Automated)

Profile Applicability:

- Level 2 - Server

Description:

The Stream Control Transmission Protocol (SCTP) is a transport layer protocol used to support message oriented communication, with several streams of messages in one connection. It serves a similar function as TCP and UDP, incorporating features of both. It is message-oriented like UDP, and ensures reliable in-sequence transport of messages with congestion control like TCP.

Rationale:

If the protocol is not being used, it is recommended that kernel module not be loaded, disabling the service to reduce the potential attack surface.

Audit:

Run the following script to verify the `sctp` module is disabled:

If the module is available in the running kernel:

- An entry including `/bin/true` or `/bin/false` exists in a file within the `/etc/modprobe.d/` directory
- The module is deny listed in a file within the `/etc/modprobe.d/` directory
- The module is not loaded in the kernel

If available in ANY installed kernel:

- The module is deny listed in a file within the `/etc/modprobe.d/` directory

If the kernel module is not available on the system, or pre-compiled into the kernel:

- No additional configuration is necessary

```

#!/usr/bin/env bash

{
    l_output="" l_output2="" l_output3="" l_dl="" # Unset output variables
    l_mname="sctp" # set module name
    l_mtype="net" # set module type
    l_searchloc="/lib/modprobe.d/*.conf /usr/local/lib/modprobe.d/*.conf /run/modprobe.d/*.conf
/etc/modprobe.d/*.conf"
    l_mpath="/lib/modules/**/kernel/$l_mtype"
    l_mpname="$(tr '-' '/' <<< "$l_mname")"
    l_mndir="$(tr '-' '/' <<< "$l_mname")"

    module_loadable_chk()
    {
        # Check if the module is currently loadable
        l_loadable="$(modprobe -n -v "$l_mname")"
        [ "$l_loadable" -gt "1" ] && l_loadable="$(grep -P --
"^(^\\h*install|\\b$l_mname)\\b" <<< "$l_loadable")"
        if grep -Pq -- '^\\h*install \\bin/(true|false)' <<< "$l_loadable"; then
            l_output="$l_output\\n - module: \"$l_mname\" is not loadable: \"$l_loadable\\n\"
        else
            l_output2="$l_output2\\n - module: \"$l_mname\" is loadable: \"$l_loadable\\n\"
        fi
    }

    module_loaded_chk()
    {
        # Check if the module is currently loaded
        if ! lsmod | grep "$l_mname" > /dev/null 2>&1; then
            l_output="$l_output\\n - module: \"$l_mname\" is not loaded"
        else
            l_output2="$l_output2\\n - module: \"$l_mname\" is loaded"
        fi
    }

    module_deny_chk()
    {
        # Check if the module is deny listed
        l_dl="y"
        if modprobe --showconfig | grep -Pq -- '^\\h*blacklist\\h+'$l_mpname'"\\b'; then
            l_output="$l_output\\n - module: \"$l_mname\" is deny listed in: \"$(grep -Pls --
"^(^\\h*blacklist\\h+$l_mname\\b" $l_searchloc)\\n\"
        else
            l_output2="$l_output2\\n - module: \"$l_mname\" is not deny listed"
        fi
    }

    # Check if the module exists on the system
    for l_mdir in $l_mpath; do
        if [ -d "$l_mdir/$l_mndir" ] && [ -n "$(ls -A $l_mdir/$l_mndir)" ]; then
            l_output3="$l_output3\\n - \"$l_mdir\\n\"
            [ "$l_dl" != "y" ] && module_deny_chk
            if [ "$l_mdir" = "/lib/modules/$(uname -r)/kernel/$l_mtype" ]; then
                module_loadable_chk
                module_loaded_chk
            fi
        else
            l_output="$l_output\\n - module: \"$l_mname\" doesn't exist in \"$l_mdir\\n\"
        fi
    done

    # Report results. If no failures output in l_output2, we pass
    [ -n "$l_output3" ] && echo -e "\\n\\n -- INFO --\\n - module: \"$l_mname\" exists in:$l_output3"
    if [ -z "$l_output2" ]; then
        echo -e "\\n- Audit Result:\\n ** PASS **\\n$l_output\\n"
    else
        echo -e "\\n- Audit Result:\\n ** FAIL **\\n - Reason(s) for audit failure:\\n$l_output2\\n"
        [ -n "$l_output" ] && echo -e "\\n- Correctly set:\\n$l_output\\n"
    fi
}

```

Remediation:

Run the following script to disable the `sctp` module:

If the module is available in the running kernel:

- Create a file with `install sctp /bin/false` in the `/etc/modprobe.d/` directory
- Create a file with `blacklist sctp` in the `/etc/modprobe.d/` directory
- Unload `sctp` from the kernel

If available in ANY installed kernel:

- Create a file with `blacklist sctp` in the `/etc/modprobe.d/` directory

If the kernel module is not available on the system or pre-compiled into the kernel:

- No remediation is necessary

```
#!/usr/bin/env bash

{
  l_mname="sctp" # set module name
  l_mtype="net" # set module type
  l_mpath="/lib/modules/**/kernel/$l_mtype"
  l_mpname="$(tr '-' '_' <<< "$l_mname")"
  l_mndir="$(tr '-' '/' <<< "$l_mname")"

  module_loadable_fix()
  {
    # If the module is currently loadable, add "install {MODULE_NAME} /bin/false" to a file in
    "/etc/modprobe.d"
    l_loadable="$(modprobe -n -v "$l_mname")"
    [ "$(wc -l <<< "$l_loadable")" -gt "1" ] && l_loadable="$(grep -P --
    "(^\\h*install|\\b$l_mname)\\b" <<< "$l_loadable")"
    if ! grep -Pq -- '^\\h*install \\bin/(true|false)' <<< "$l_loadable"; then
      echo -e "\\n - setting module: \"$l_mname\" to be not loadable"
      echo -e "install $l_mname /bin/false" >> /etc/modprobe.d/"$l_mpname".conf
    fi
  }
  module_loaded_fix()
  {
    # If the module is currently loaded, unload the module
    if lsmod | grep "$l_mname" > /dev/null 2>&1; then
      echo -e "\\n - unloading module \"$l_mname\""
      modprobe -r "$l_mname"
    fi
  }
  module_deny_fix()
  {
    # If the module isn't deny listed, denylist the module
    if ! modprobe --showconfig | grep -Pq -- "^\\h*blacklist\\h+$l_mpname\\b"; then
      echo -e "\\n - deny listing \"$l_mname\""
      echo -e "blacklist $l_mname" >> /etc/modprobe.d/"$l_mpname".conf
    fi
  }
  # Check if the module exists on the system
  for l_mdir in $l_mpath; do
    if [ -d "$l_mdir/$l_mndir" ] && [ -n "$(ls -A $l_mdir/$l_mndir)" ]; then
      echo -e "\\n - module: \"$l_mname\" exists in \"$l_mdir\"\\n - checking if disabled..."
      module_deny_fix
      if [ "$l_mdir" = "/lib/modules/$(uname -r)/kernel/$l_mtype" ]; then
        module_loadable_fix
        module_loaded_fix
      fi
    else
      echo -e "\\n - module: \"$l_mname\" doesn't exist in \"$l_mdir\"\\n"
    fi
  done
  echo -e "\\n - remediation of module: \"$l_mname\" complete\\n"
}

```

References:

1. NIST SP 800-53 Rev. 5: SI-4, CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1068, T1068.000, T1210, T1210.000	TA0008	M1042

3.1.4 Ensure RDS is disabled (Automated)

Profile Applicability:

- Level 2 - Server

Description:

The Reliable Datagram Sockets (RDS) protocol is a transport layer protocol designed to provide low-latency, high-bandwidth communications between cluster nodes. It was developed by the Oracle Corporation.

Rationale:

If the protocol is not being used, it is recommended that kernel module not be loaded, disabling the service to reduce the potential attack surface.

Audit:

Run the following script to verify the `rds` module is disabled:

If the module is available in the running kernel:

- An entry including `/bin/true` or `/bin/false` exists in a file within the `/etc/modprobe.d/` directory
- The module is deny listed in a file within the `/etc/modprobe.d/` directory
- The module is not loaded in the kernel

If available in ANY installed kernel:

- The module is deny listed in a file within the `/etc/modprobe.d/` directory

If the kernel module is not available on the system, or pre-compiled into the kernel:

- No additional configuration is necessary

```

#!/usr/bin/env bash

{
    l_output="" l_output2="" l_output3="" l_dl="" # Unset output variables
    l_mname="rds" # set module name
    l_mtype="net" # set module type
    l_searchloc="/lib/modprobe.d/*.conf /usr/local/lib/modprobe.d/*.conf /run/modprobe.d/*.conf
/etc/modprobe.d/*.conf"
    l_mpath="/lib/modules/**/kernel/$l_mtype"
    l_mpname="$(tr '-' '/' <<< "$l_mname")"
    l_mndir="$(tr '-' '/' <<< "$l_mname")"

    module_loadable_chk()
    {
        # Check if the module is currently loadable
        l_loadable="$(modprobe -n -v "$l_mname")"
        [ "$l_loadable" -gt "1" ] && l_loadable="$(grep -P --
"^(h*install|b$l_mname)b" <<< "$l_loadable")"
        if grep -Pq -- '^(h*install |bin|/true|false)' <<< "$l_loadable"; then
            l_output="$l_output\n - module: \"$l_mname\" is not loadable: \"$l_loadable\""
        else
            l_output2="$l_output2\n - module: \"$l_mname\" is loadable: \"$l_loadable\""
        fi
    }

    module_loaded_chk()
    {
        # Check if the module is currently loaded
        if ! lsmod | grep "$l_mname" > /dev/null 2>&1; then
            l_output="$l_output\n - module: \"$l_mname\" is not loaded"
        else
            l_output2="$l_output2\n - module: \"$l_mname\" is loaded"
        fi
    }

    module_deny_chk()
    {
        # Check if the module is deny listed
        l_dl="y"
        if modprobe --showconfig | grep -Pq -- '^(h*blacklist|h+"$l_mpname"'b'; then
            l_output="$l_output\n - module: \"$l_mname\" is deny listed in: \"$(grep -Pls --
"^(h*blacklist|h+$l_mname)b" $l_searchloc)\""
        else
            l_output2="$l_output2\n - module: \"$l_mname\" is not deny listed"
        fi
    }

    # Check if the module exists on the system
    for l_mdir in $l_mpath; do
        if [ -d "$l_mdir/$l_mndir" ] && [ -n "$(ls -A $l_mdir/$l_mndir)" ]; then
            l_output3="$l_output3\n - \"$l_mdir\""
            [ "$l_dl" != "y" ] && module_deny_chk
            if [ "$l_mdir" = "/lib/modules/$(uname -r)/kernel/$l_mtype" ]; then
                module_loadable_chk
                module_loaded_chk
            fi
        else
            l_output="$l_output\n - module: \"$l_mname\" doesn't exist in \"$l_mdir\""
        fi
    done

    # Report results. If no failures output in l_output2, we pass
    [ -n "$l_output3" ] && echo -e "\n\n -- INFO --\n - module: \"$l_mname\" exists in:$l_output3"
    if [ -z "$l_output2" ]; then
        echo -e "\n- Audit Result:\n ** PASS **\n$l_output\n"
    else
        echo -e "\n- Audit Result:\n ** FAIL **\n - Reason(s) for audit failure:\n$l_output2\n"
        [ -n "$l_output" ] && echo -e "\n- Correctly set:\n$l_output\n"
    fi
}

```

Remediation:

Run the following script to disable the `rds` module:

If the module is available in the running kernel:

- Create a file with `install rds /bin/false` in the `/etc/modprobe.d/` directory
- Create a file with `blacklist rds` in the `/etc/modprobe.d/` directory
- Unload `rds` from the kernel

If available in ANY installed kernel:

- Create a file with `blacklist rds` in the `/etc/modprobe.d/` directory

If the kernel module is not available on the system or pre-compiled into the kernel:

- No remediation is necessary

```

#!/usr/bin/env bash

{
  l_mname="rds" # set module name
  l_mtype="net" # set module type
  l_mpath="/lib/modules/**/kernel/$l_mtype"
  l_mpname="$(tr '-' '_' <<< "$l_mname")"
  l_mndir="$(tr '-' '/' <<< "$l_mname")"

  module_loadable_fix()
  {
    # If the module is currently loadable, add "install {MODULE_NAME} /bin/false" to a file in
    "/etc/modprobe.d"
    l_loadable="$(modprobe -n -v "$l_mname")"
    [ "$(wc -l <<< "$l_loadable")" -gt "1" ] && l_loadable="$(grep -P --
    "(^\\h*install|\\b$l_mname)\\b" <<< "$l_loadable")"
    if ! grep -Pq -- '^\\h*install \\bin/(true|false)' <<< "$l_loadable"; then
      echo -e "\\n - setting module: \"$l_mname\" to be not loadable"
      echo -e "install $l_mname /bin/false" >> /etc/modprobe.d/"$l_mpname".conf
    fi
  }

  module_loaded_fix()
  {
    # If the module is currently loaded, unload the module
    if lsmod | grep "$l_mname" > /dev/null 2>&1; then
      echo -e "\\n - unloading module \"$l_mname\""
      modprobe -r "$l_mname"
    fi
  }

  module_deny_fix()
  {
    # If the module isn't deny listed, denylist the module
    if ! modprobe --showconfig | grep -Pq -- "^\\h*blacklist\\h+$l_mpname\\b"; then
      echo -e "\\n - deny listing \"$l_mname\""
      echo -e "blacklist $l_mname" >> /etc/modprobe.d/"$l_mpname".conf
    fi
  }

  # Check if the module exists on the system
  for l_mdir in $l_mpath; do
    if [ -d "$l_mdir/$l_mndir" ] && [ -n "$(ls -A $l_mdir/$l_mndir)" ]; then
      echo -e "\\n - module: \"$l_mname\" exists in \"$l_mdir\"\\n - checking if disabled..."
      module_deny_fix
      if [ "$l_mdir" = "/lib/modules/$(uname -r)/kernel/$l_mtype" ]; then
        module_loadable_fix
        module_loaded_fix
      fi
    else
      echo -e "\\n - module: \"$l_mname\" doesn't exist in \"$l_mdir\"\\n"
    fi
  done
  echo -e "\\n - remediation of module: \"$l_mname\" complete\\n"
}

```

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1068, T1068.000, T1210, T1210.000	TA0008	M1042

3.1.5 Ensure TIPC is disabled (Automated)

Profile Applicability:

- Level 2 - Server

Description:

The Transparent Inter-Process Communication (TIPC) protocol is designed to provide communication between cluster nodes.

Rationale:

If the protocol is not being used, it is recommended that kernel module not be loaded, disabling the service to reduce the potential attack surface.

Audit:

Run the following script to verify the `tipc` module is disabled:

If the module is available in the running kernel:

- An entry including `/bin/true` or `/bin/false` exists in a file within the `/etc/modprobe.d/` directory
- The module is deny listed in a file within the `/etc/modprobe.d/` directory
- The module is not loaded in the kernel

If available in ANY installed kernel:

- The module is deny listed in a file within the `/etc/modprobe.d/` directory

If the kernel module is not available on the system, or pre-compiled into the kernel:

- No additional configuration is necessary

```

#!/usr/bin/env bash

{
    l_output="" l_output2="" l_output3="" l_dl="" # Unset output variables
    l_mname="tipc" # set module name
    l_mtype="net" # set module type
    l_searchloc="/lib/modprobe.d/*.conf /usr/local/lib/modprobe.d/*.conf /run/modprobe.d/*.conf
/etc/modprobe.d/*.conf"
    l_mpath="/lib/modules/**/kernel/$l_mtype"
    l_mpname="$(tr '-' '/' <<< "$l_mname")"
    l_mndir="$(tr '-' '/' <<< "$l_mname")"

    module_loadable_chk()
    {
        # Check if the module is currently loadable
        l_loadable="$(modprobe -n -v "$l_mname")"
        [ "$l_loadable" -gt "1" ] && l_loadable="$(grep -P --
"^(^\\h*install|\\b$l_mname)\\b" <<< "$l_loadable")"
        if grep -Pq -- '^\\h*install \\bin/(true|false)' <<< "$l_loadable"; then
            l_output="$l_output\\n - module: \"$l_mname\" is not loadable: \"$l_loadable\\n\"
        else
            l_output2="$l_output2\\n - module: \"$l_mname\" is loadable: \"$l_loadable\\n\"
        fi
    }

    module_loaded_chk()
    {
        # Check if the module is currently loaded
        if ! lsmod | grep "$l_mname" > /dev/null 2>&1; then
            l_output="$l_output\\n - module: \"$l_mname\" is not loaded"
        else
            l_output2="$l_output2\\n - module: \"$l_mname\" is loaded"
        fi
    }

    module_deny_chk()
    {
        # Check if the module is deny listed
        l_dl="y"
        if modprobe --showconfig | grep -Pq -- '^\\h*blacklist\\h+'$l_mpname'"\\b'; then
            l_output="$l_output\\n - module: \"$l_mname\" is deny listed in: \"$(grep -Pls --
"^(^\\h*blacklist\\h+$l_mname\\b" $l_searchloc)\\n\"
        else
            l_output2="$l_output2\\n - module: \"$l_mname\" is not deny listed"
        fi
    }

    # Check if the module exists on the system
    for l_mdir in $l_mpath; do
        if [ -d "$l_mdir/$l_mndir" ] && [ -n "$(ls -A $l_mdir/$l_mndir)" ]; then
            l_output3="$l_output3\\n - \"$l_mdir\\n\"
            [ "$l_dl" != "y" ] && module_deny_chk
            if [ "$l_mdir" = "/lib/modules/$(uname -r)/kernel/$l_mtype" ]; then
                module_loadable_chk
                module_loaded_chk
            fi
        else
            l_output="$l_output\\n - module: \"$l_mname\" doesn't exist in \"$l_mdir\\n\"
        fi
    done

    # Report results. If no failures output in l_output2, we pass
    [ -n "$l_output3" ] && echo -e "\\n\\n -- INFO --\\n - module: \"$l_mname\" exists in:$l_output3"
    if [ -z "$l_output2" ]; then
        echo -e "\\n- Audit Result:\\n ** PASS **\\n$l_output\\n"
    else
        echo -e "\\n- Audit Result:\\n ** FAIL **\\n - Reason(s) for audit failure:\\n$l_output2\\n"
        [ -n "$l_output" ] && echo -e "\\n- Correctly set:\\n$l_output\\n"
    fi
}

```


Remediation:

Run the following script to disable the `tipc` module:

If the module is available in the running kernel:

- Create a file with `install tipc /bin/false` in the `/etc/modprobe.d/` directory
- Create a file with `blacklist tipc` in the `/etc/modprobe.d/` directory
- Unload `tipc` from the kernel

If available in ANY installed kernel:

- Create a file with `blacklist tipc` in the `/etc/modprobe.d/` directory

If the kernel module is not available on the system or pre-compiled into the kernel:

- No remediation is necessary

```
#!/usr/bin/env bash

{
  l_mname="tipc" # set module name
  l_mtype="net" # set module type
  l_mpath="/lib/modules/**/kernel/$l_mtype"
  l_mpname="$(tr '-' '_' <<< "$l_mname")"
  l_mndir="$(tr '-' '/' <<< "$l_mname")"

  module_loadable_fix()
  {
    # If the module is currently loadable, add "install {MODULE_NAME} /bin/false" to a file in
    "/etc/modprobe.d"
    l_loadable="$(modprobe -n -v "$l_mname")"
    [ "$(wc -l <<< "$l_loadable")" -gt "1" ] && l_loadable="$(grep -P --
    "(^\\h*install|\\b$l_mname)\\b" <<< "$l_loadable")"
    if ! grep -Pq -- '^\\h*install \\bin/(true|false)' <<< "$l_loadable"; then
      echo -e "\\n - setting module: \"$l_mname\" to be not loadable"
      echo -e "install $l_mname /bin/false" >> /etc/modprobe.d/"$l_mpname".conf
    fi
  }

  module_loaded_fix()
  {
    # If the module is currently loaded, unload the module
    if lsmod | grep "$l_mname" > /dev/null 2>&1; then
      echo -e "\\n - unloading module \"$l_mname\""
      modprobe -r "$l_mname"
    fi
  }

  module_deny_fix()
  {
    # If the module isn't deny listed, denylist the module
    if ! modprobe --showconfig | grep -Pq -- "^\\h*blacklist\\h+$l_mpname\\b"; then
      echo -e "\\n - deny listing \"$l_mname\""
      echo -e "blacklist $l_mname" >> /etc/modprobe.d/"$l_mpname".conf
    fi
  }

  # Check if the module exists on the system
  for l_mdir in $l_mpath; do
    if [ -d "$l_mdir/$l_mndir" ] && [ -n "$(ls -A $l_mdir/$l_mndir)" ]; then
      echo -e "\\n - module: \"$l_mname\" exists in \"$l_mdir\"\\n - checking if disabled..."
      module_deny_fix
      if [ "$l_mdir" = "/lib/modules/$(uname -r)/kernel/$l_mtype" ]; then
        module_loadable_fix
        module_loaded_fix
      fi
    else
      echo -e "\\n - module: \"$l_mname\" doesn't exist in \"$l_mdir\"\\n"
    fi
  done
  echo -e "\\n - remediation of module: \"$l_mname\" complete\\n"
}

```

References:

1. NIST SP 800-53 Rev. 5: SI-4, CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1068, T1068.000, T1210, T1210.000	TA0008	M1042

3.2 Configure Network Parameters (Host Only)

The following network parameters are intended for use if the system is to act as a host only. A system is considered host only if the system has a single interface, or has multiple interfaces but will not be configured as a router.

Note:

- `sysctl` settings are defined through files in `/usr/local/lib`, `/usr/lib/`, `/lib/`, `/run/`, and `/etc/`
- Files are typically placed in the `sysctl.d` directory within the parent directory
- The paths where `sysctl` preload files usually exist
 - `/run/sysctl.d/*.conf`
 - `/etc/sysctl.d/*.conf`
 - `/usr/local/lib/sysctl.d/*.conf`
 - `/usr/lib/sysctl.d/*.conf`
 - `/lib/sysctl.d/*.conf`
 - `/etc/sysctl.conf`
- Files must have the `".conf"` extension
- Vendors settings usually live in `/usr/lib/` or `/usr/local/lib/`
- To override a whole file, create a new file with the same name in `/etc/sysctl.d/` and put new settings there.
- To override only specific settings, add a file with a lexically later name in `/etc/sysctl.d/` and put new settings there.
- The command `/usr/lib/systemd/systemd-sysctl --cat-config` produces output containing The system's loaded kernel parameters and the files they're configured in:
 - Entries listed latter in the file take precedence over the same settings listed earlier in the file
 - Files containing kernel parameters that are over-ridden by other files with the same name will not be listed
 - On systems running UncomplicatedFirewall, the kernel parameters may be set or over-written. This will not be visible in the output of the command

The system's loaded kernel parameters and the files they're configured in can be viewed by running the following command:

```
# /usr/lib/systemd/systemd-sysctl --cat-config
```

3.2.1 Ensure IP forwarding is disabled (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `net.ipv4.ip_forward` and `net.ipv6.conf.all.forwarding` flags are used to tell the system whether it can forward packets or not.

Rationale:

Setting `net.ipv4.ip_forward` and `net.ipv6.conf.all.forwarding` to 0 ensures that a system with multiple interfaces (for example, a hard proxy), will never be able to forward packets, and therefore, never serve as a router.

Audit:

Run the following script to verify the following kernel parameters are set in the running configuration and correctly loaded from a kernel parameter configuration file:

- `net.ipv4.ip_forward` is set to 0
- `net.ipv6.conf.all.forwarding` is set to 0

Note:

- kernel parameters are loaded by file and parameter order precedence. The following script observes this precedence as part of the auditing procedure. The parameters being checked may be set correctly in a file. If that file is superseded, the parameter is overridden by an incorrect setting later in that file, or in a conically later file, that "correct" setting will be ignored both by the script and by the system during a normal kernel parameter load sequence.
- IPv6 kernel parameters only apply to systems where IPv6 is enabled

```

#!/usr/bin/env bash

{
    l_output="" l_output2=""
    a_parlist=("net.ipv4.ip_forward=0" "net.ipv6.conf.all.forwarding=0")
    l_ufwscf="$([ -f /etc/default/ufw ] && awk -F= '/^\s*IPT_SYSCTL=/ {print $2}'
/etc/default/ufw)"
    kernel_parameter_chk()
    {
        l_krp="$(sysctl "$l_kpname" | awk -F= '{print $2}' | xargs)" # Check running configuration
        if [ "$l_krp" = "$l_kpvalue" ]; then
            l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_krp\" in the running
configuration"
        else
            l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_krp\" in the running
configuration and should have a value of: \"$l_kpvalue\""
        fi
        unset A_out; declare -A A_out # Check durable setting (files)
        while read -r l_out; do
            if [ -n "$l_out" ]; then
                if [[ $l_out =~ ^\s*# ]]; then
                    l_file="$l_out// # /)"
                else
                    l_kpar="$(awk -F= '{print $1}' <<< "$l_out" | xargs)"
                    [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]=" $l_file")
                fi
            fi
            done < <(/usr/lib/systemd/systemd-sysctl --cat-config | grep -Po
'^\h*([^\#\n\r]+|#\h*\[/^\#\n\r\h]+\.\conf\b)')
            if [ -n "$l_ufwscf" ]; then # Account for systems with UFW (Not covered by systemd-sysctl -
-cat-config)
                l_kpar="$(grep -Po '^ \h*$l_kpname\b' "$l_ufwscf" | xargs)"
                l_kpar="$l_kpar//\./)"
                [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]=" $l_ufwscf")
            fi
            if (( ${#A_out[@]} > 0 )); then # Assess output from files and generate output
                while IFS="" read -r l_fkpname l_fkpvalue; do
                    l_fkpname="$l_fkpname// /)"; l_fkpvalue="$l_fkpvalue// /)"
                    if [ "$l_fkpvalue" = "$l_kpvalue" ]; then
                        l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_fkpvalue\" in
\"$(printf '%s' "${A_out[@]}")\""\n"
                    else
                        l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_fkpvalue\" in
\"$(printf '%s' "${A_out[@]}")\"" and should have a value of: \"$l_kpvalue\""\n"
                    fi
                    done < <(grep -Po -- "\^ \h*$l_kpname\h*=\h*\H+" "${A_out[@]}")
                else
                    l_output2="$l_output2\n - \"$l_kpname\" is not set in an included file\n    ** Note:
\"$l_kpname\" May be set in a file that's ignored by load procedure **\n"
                fi
            }
            while IFS="" read -r l_kpname l_kpvalue; do # Assess and check parameters
                l_kpname="$l_kpname// /)"; l_kpvalue="$l_kpvalue// /)"
                if ! grep -Pqs '\^ \h*0\b' /sys/module/ipv6/parameters/disable && grep -q '\^net.ipv6.' <<<
"$l_kpname"; then
                    l_output="$l_output\n - IPv6 is disabled on the system, \"$l_kpname\" is not applicable"
                else
                    kernel_parameter_chk
                fi
                done < <(printf '%s\n' "${a_parlist[@]}")
                if [ -z "$l_output2" ]; then # Provide output from checks
                    echo -e "\n- Audit Result:\n    ** PASS **\n$l_output\n"
                else
                    echo -e "\n- Audit Result:\n    ** FAIL **\n - Reason(s) for audit failure:\n$l_output2\n"
                    [ -n "$l_output" ] && echo -e "\n- Correctly set:\n$l_output\n"
                fi
            fi
        }
    }
}

```

Remediation:

Set the following parameter in `/etc/sysctl.conf` or a file in `/etc/sysctl.d/` ending in `.conf`:

- `net.ipv4.ip_forward = 0`

Example:

```
# printf "  
net.ipv4.ip_forward = 0  
" >> /etc/sysctl.d/60-netipv4_sysctl.conf
```

Run the following command to set the active kernel parameters:

```
# {  
    sysctl -w net.ipv4.ip_forward=0  
    sysctl -w net.ipv4.route.flush=1  
}
```

-IF- IPv6 is enabled on the system:

Set the following parameter in `/etc/sysctl.conf` or a file in `/etc/sysctl.d/` ending in `.conf`:

- `net.ipv6.conf.all.forwarding = 0`

Example:

```
# printf "  
net.ipv6.conf.all.forwarding = 0  
" >> /etc/sysctl.d/60-netipv6_sysctl.conf
```

Run the following command to set the active kernel parameters:

```
# {  
    sysctl -w net.ipv6.conf.all.forwarding=0  
    sysctl -w net.ipv6.route.flush=1  
}
```

Note: If these settings appear in a conically later file, or later in the same file, these settings will be overwritten

Default Value:

`net.ipv4.ip_forward = 0`

`net.ipv6.conf.all.forwarding = 0`

References:

1. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1557, T1557.000	TA0006, TA0009	M1030, M1042

3.2.2 *Ensure packet redirect sending is disabled (Automated)*

Profile Applicability:

- Level 1 - Server

Description:

ICMP Redirects are used to send routing information to other hosts. As a host itself does not act as a router (in a host only configuration), there is no need to send redirects.

Rationale:

An attacker could use a compromised host to send invalid ICMP redirects to other router devices in an attempt to corrupt routing and have users access a system set up by the attacker as opposed to a valid system.

Audit:

Run the following script to verify the following kernel parameters are set in the running configuration and correctly loaded from a kernel parameter configuration file:

- `net.ipv4.conf.all.send_redirects` is set to 0
- `net.ipv4.conf.default.send_redirects` is set to 0

Note: kernel parameters are loaded by file and parameter order precedence. The following script observes this precedence as part of the auditing procedure. The parameters being checked may be set correctly in a file. If that file is superseded, the parameter is overridden by an incorrect setting later in that file, or in a conically later file, that "correct" setting will be ignored both by the script and by the system during a normal kernel parameter load sequence.

```
#!/usr/bin/env bash

{
    l_output="" l_output2=""
    a_parlist=("net.ipv4.conf.all.send_redirects=0" "net.ipv4.conf.default.send_redirects=0")
    l_ufwscf="$( [ -f /etc/default/ufw ] && awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw )"
    kernel_parameter_chk()
    {
        l_krp="$(sysctl "$l_kpname" | awk -F= '{print $2}' | xargs)" # Check running configuration
        if [ "$l_krp" = "$l_kpvalue" ]; then
            l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_krp\" in the running configuration"
        else
            l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_krp\" in the running configuration and should have a value of: \"$l_kpvalue\""
        fi
        unset A_out; declare -A A_out # Check durable setting (files)
        while read -r l_out; do
            if [ -n "$l_out" ]; then
                if [[ $l_out =~ ^\s*# ]]; then
                    l_file="$l_out// # /)"
                else
                    l_kpar="$(awk -F= '{print $1}' <<< "$l_out" | xargs)"
                    [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]=" $l_file")
                fi
            fi
            done < <(/usr/lib/systemd/systemd-sysctl --cat-config | grep -Po
'^\h*([^\#\n\r]+|#\h*\[/^\#\n\r\h]+\.\conf\b)')
            if [ -n "$l_ufwscf" ]; then # Account for systems with UFW (Not covered by systemd-sysctl -cat-config)
                l_kpar="$(grep -Po '^h*$l_kpname\b' "$l_ufwscf" | xargs)"
                l_kpar="$l_kpar//\./)"
                [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]=" $l_ufwscf")
            fi
            if (( ${#A_out[@]} > 0 )); then # Assess output from files and generate output
                while IFS="" read -r l_fkpname l_fkpvalue; do
                    l_fkpname="$l_fkpname// /)"; l_fkpvalue="$l_fkpvalue// /)"
                    if [ "$l_fkpvalue" = "$l_kpvalue" ]; then
                        l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_krp\" in \"$(printf '%s' "${A_out[@]}")\""\n"
                    else
                        l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_fkpvalue\" in \"$(printf '%s' "${A_out[@]}")\" and should have a value of: \"$l_kpvalue\""\n"
                    fi
                    done < <(grep -Po -- "\^h*$l_kpname\h*=\h*\H+" "${A_out[@]}")
                else
                    l_output2="$l_output2\n - \"$l_kpname\" is not set in an included file\n    ** Note: \"$l_kpname\" May be set in a file that's ignored by load procedure **\n"
                fi
            }
            while IFS="" read -r l_kpname l_kpvalue; do # Assess and check parameters
                l_kpname="$l_kpname// /)"; l_kpvalue="$l_kpvalue// /)"
                if ! grep -Pqs '\^h*0\b' /sys/module/ipv6/parameters/disable && grep -q '\^net.ipv6.' <<< "$l_kpname"; then
                    l_output="$l_output\n - IPv6 is disabled on the system, \"$l_kpname\" is not applicable"
                else
                    kernel_parameter_chk
                fi
                done < <(printf '%s\n' "${a_parlist[@]}")
                if [ -z "$l_output2" ]; then # Provide output from checks
                    echo -e "\n- Audit Result:\n    ** PASS **\n$l_output\n"
                else
                    echo -e "\n- Audit Result:\n    ** FAIL **\n - Reason(s) for audit failure:\n$l_output2\n"
                    [ -n "$l_output" ] && echo -e "\n- Correctly set:\n$l_output\n"
                fi
            fi
        }
    }
}
```

Remediation:

Set the following parameters in `/etc/sysctl.conf` or a file in `/etc/sysctl.d/` ending in `.conf`:

- `net.ipv4.conf.all.send_redirects = 0`
- `net.ipv4.conf.default.send_redirects = 0`

Example:

```
# printf "  
net.ipv4.conf.all.send_redirects = 0  
net.ipv4.conf.default.send_redirects = 0  
" >> /etc/sysctl.d/60-netipv4_sysctl.conf
```

Run the following command to set the active kernel parameters:

```
# {  
  sysctl -w net.ipv4.conf.all.send_redirects=0  
  sysctl -w net.ipv4.conf.default.send_redirects=0  
  sysctl -w net.ipv4.route.flush=1  
}
```

Note: If these settings appear in a conically later file, or later in the same file, these settings will be overwritten

References:

1. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 <u>Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●
v7	9.2 <u>Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.		●	●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1557, T1557.000	TA0006, TA0009	M1030, M1042

3.3 Configure Network Parameters (Host and Router)

The following network parameters are intended for use on both host only and router systems. A system acts as a router if it has at least two interfaces and is configured to perform routing functions.

Note:

- `sysctl` settings are defined through files in `/usr/local/lib`, `/usr/lib/`, `/lib/`, `/run/`, and `/etc/`
- Files are typically placed in the `sysctl.d` directory within the parent directory
- The paths where `sysctl` preload files usually exist
 - `/run/sysctl.d/*.conf`
 - `/etc/sysctl.d/*.conf`
 - `/usr/local/lib/sysctl.d/*.conf`
 - `/usr/lib/sysctl.d/*.conf`
 - `/lib/sysctl.d/*.conf`
 - `/etc/sysctl.conf`
- Files must have the `".conf"` extension
- Vendors settings usually live in `/usr/lib/` or `/usr/local/lib/`
- To override a whole file, create a new file with the same name in `/etc/sysctl.d/` and put new settings there.
- To override only specific settings, add a file with a lexically later name in `/etc/sysctl.d/` and put new settings there.
- The command `/usr/lib/systemd/systemd-sysctl --cat-config` produces output containing The system's loaded kernel parameters and the files they're configured in:
 - Entries listed latter in the file take precedence over the same settings listed earlier in the file
 - Files containing kernel parameters that are over-ridden by other files with the same name will not be listed
 - On systems running UncomplicatedFirewall, the kernel parameters may be set or over-written. This will not be visible in the output of the command

The system's loaded kernel parameters and the files they're configured in can be viewed by running the following command:

```
# /usr/lib/systemd/systemd-sysctl --cat-config
```

3.3.1 Ensure source routed packets are not accepted (Automated)

Profile Applicability:

- Level 1 - Server

Description:

In networking, source routing allows a sender to partially or fully specify the route packets take through a network. In contrast, non-source routed packets travel a path determined by routers in the network. In some cases, systems may not be routable or reachable from some locations (e.g. private addresses vs. Internet routable), and so source routed packets would need to be used.

Rationale:

Setting `net.ipv4.conf.all.accept_source_route`, `net.ipv4.conf.default.accept_source_route`, `net.ipv6.conf.all.accept_source_route` and `net.ipv6.conf.default.accept_source_route` to 0 disables the system from accepting source routed packets. Assume this system was capable of routing packets to Internet routable addresses on one interface and private addresses on another interface. Assume that the private addresses were not routable to the Internet routable addresses and vice versa. Under normal routing circumstances, an attacker from the Internet routable addresses could not use the system as a way to reach the private address systems. If, however, source routed packets were allowed, they could be used to gain access to the private address systems as the route could be specified, rather than rely on routing protocols that did not allow this routing.

Audit:

Run the following script to verify the following kernel parameters are set in the running configuration and correctly loaded from a kernel parameter configuration file:

- `net.ipv4.conf.all.accept_source_route` is set to 0
- `net.ipv4.conf.default.accept_source_route` is set to 0
- `net.ipv6.conf.all.accept_source_route` is set to 0
- `net.ipv6.conf.default.accept_source_route` is set to 0

Note:

- kernel parameters are loaded by file and parameter order precedence. The following script observes this precedence as part of the auditing procedure. The parameters being checked may be set correctly in a file. If that file is superseded, the parameter is overridden by an incorrect setting later in that file, or in a conically later file, that "correct" setting will be ignored both by the script and by the system during a normal kernel parameter load sequence.
- IPv6 kernel parameters only apply to systems where IPv6 is enabled

```
#!/usr/bin/env bash

{
    l_output="" l_output2=""
    a_parlist=("net.ipv4.conf.all.accept_source_route=0"
"net.ipv4.conf.default.accept_source_route=0" "net.ipv6.conf.all.accept_source_route=0"
"net.ipv6.conf.default.accept_source_route=0")
    l_ufwscf="$([ -f /etc/default/ufw ] && awk -F= '/^\s*IPT_SYSCTL=/ {print $2}'
/etc/default/ufw)"
    kernel_parameter_chk()
    {
        l_krp="$(sysctl "$l_kpname" | awk -F= '{print $2}' | xargs)" # Check running configuration
        if [ "$l_krp" = "$l_kpvalue" ]; then
            l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_krp\" in the running
configuration"
        else
            l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_krp\" in the running
configuration and should have a value of: \"$l_kpvalue\""
        fi
        unset A_out; declare -A A_out # Check durable setting (files)
        while read -r l_out; do
            if [ -n "$l_out" ]; then
                if [[ $l_out =~ ^\s*# ]]; then
                    l_file="${l_out//\# /}"
                else
                    l_kpar="$(awk -F= '{print $1}' <<< "$l_out" | xargs)"
                    [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]="$l_file")
                fi
            fi
        done < <(/usr/lib/systemd/systemd-sysctl --cat-config | grep -Po
'^\h*([^\n\r]+|#\h*\[/^\n\r\h]+\.\conf\b)')
        if [ -n "$l_ufwscf" ]; then # Account for systems with UFW (Not covered by systemd-sysctl -
-cat-config)
            l_kpar="$(grep -Po "^\h*$l_kpname\b" "$l_ufwscf" | xargs)"
            l_kpar="${l_kpar//\./}"
            [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]="$l_ufwscf")
        fi
        if (( ${#A_out[@]} > 0 )); then # Assess output from files and generate output
            while IFS="" read -r l_fkpname l_fkpvalue; do
                l_fkpname="${l_fkpname// /}"; l_fkpvalue="${l_fkpvalue// /}"
                if [ "$l_fkpvalue" = "$l_kpvalue" ]; then
                    l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_fkpvalue\" in
\"$(printf '%s' "${A_out[@]}")\""\n"
                else
                    l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_fkpvalue\" in
\"$(printf '%s' "${A_out[@]}")\" and should have a value of: \"$l_kpvalue\""\n"
                fi
                done < <(grep -Po -- "^\h*$l_kpname\h*=\h*\H+" "${A_out[@]}")
            else
                l_output2="$l_output2\n - \"$l_kpname\" is not set in an included file\n    ** Note:
\"$l_kpname\" May be set in a file that's ignored by load procedure **\n"
            fi
        }
        while IFS="" read -r l_kpname l_kpvalue; do # Assess and check parameters
            l_kpname="${l_kpname// /}"; l_kpvalue="${l_kpvalue// /}"
            if ! grep -Pqs '^\h*0\b' /sys/module/ipv6/parameters/disable && grep -q 'net.ipv6.' <<<
"$l_kpname"; then
                l_output="$l_output\n - IPv6 is disabled on the system, \"$l_kpname\" is not applicable"
            else
                kernel_parameter_chk
            fi
        done < <(printf '%s\n' "${a_parlist[@]}")
        if [ -z "$l_output2" ]; then # Provide output from checks
            echo -e "\n- Audit Result:\n    ** PASS **\n$l_output\n"
        else
            echo -e "\n- Audit Result:\n    ** FAIL **\n - Reason(s) for audit failure:\n$l_output2\n"
            [ -n "$l_output" ] && echo -e "\n- Correctly set:\n$l_output\n"
        fi
    }
}
```


Remediation:

Set the following parameters in `/etc/sysctl.conf` or a file in `/etc/sysctl.d/` ending in `.conf`:

- `net.ipv4.conf.all.accept_source_route = 0`
- `net.ipv4.conf.default.accept_source_route = 0`

Example:

```
# printf "  
net.ipv4.conf.all.accept_source_route = 0  
net.ipv4.conf.default.accept_source_route = 0  
" >> /etc/sysctl.d/60-netipv4_sysctl.conf
```

Run the following command to set the active kernel parameters:

```
# {  
    sysctl -w net.ipv4.conf.all.accept_source_route=0  
    sysctl -w net.ipv4.conf.default.accept_source_route=0  
    sysctl -w net.ipv4.route.flush=1  
}
```

-IF- IPv6 is enabled on the system:

Set the following parameters in `/etc/sysctl.conf` or a file in `/etc/sysctl.d/` ending in `.conf`:

- `net.ipv6.conf.all.accept_source_route = 0`
- `net.ipv6.conf.default.accept_source_route = 0`

Example:

```
# printf "  
net.ipv6.conf.all.accept_source_route = 0  
net.ipv6.conf.default.accept_source_route = 0  
" >> /etc/sysctl.d/60-netipv6_sysctl.conf
```

Run the following command to set the active kernel parameters:

```
# {  
    sysctl -w net.ipv6.conf.all.accept_source_route=0  
    sysctl -w net.ipv6.conf.default.accept_source_route=0  
    sysctl -w net.ipv6.route.flush=1  
}
```

Note: If these settings appear in a conically later file, or later in the same file, these settings will be overwritten

Default Value:

`net.ipv4.conf.all.accept_source_route = 0`

`net.ipv4.conf.default.accept_source_route = 0`

`net.ipv6.conf.all.accept_source_route = 0`

`net.ipv6.conf.default.accept_source_route = 0`

References:

1. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 <u>Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	9.2 <u>Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1590, T1590.005	TA0007	

3.3.2 Ensure ICMP redirects are not accepted (Automated)

Profile Applicability:

- Level 1 - Server

Description:

ICMP redirect messages are packets that convey routing information and tell your host (acting as a router) to send packets via an alternate path. It is a way of allowing an outside routing device to update your system routing tables.

Rationale:

ICMP redirect messages are packets that convey routing information and tell your host (acting as a router) to send packets via an alternate path. It is a way of allowing an outside routing device to update your system routing tables. By setting

`net.ipv4.conf.all.accept_redirects`, `net.ipv4.conf.default.accept_redirects`, `net.ipv6.conf.all.accept_redirects`, and `net.ipv6.conf.default.accept_redirects` to 0, the system will not accept any ICMP redirect messages, and therefore, won't allow outsiders to update the system's routing tables.

Audit:

Run the following script to verify the following kernel parameters are set in the running configuration and correctly loaded from a kernel parameter configuration file:

- `net.ipv4.conf.all.accept_redirects` is set to 0
- `net.ipv4.conf.default.accept_redirects` is set to 0
- `net.ipv6.conf.all.accept_redirects` is set to 0
- `net.ipv6.conf.default.accept_redirects` is set to 0

Note:

- kernel parameters are loaded by file and parameter order precedence. The following script observes this precedence as part of the auditing procedure. The parameters being checked may be set correctly in a file. If that file is superseded, the parameter is overridden by an incorrect setting later in that file, or in a conically later file, that "correct" setting will be ignored both by the script and by the system during a normal kernel parameter load sequence.
- IPv6 kernel parameters only apply to systems where IPv6 is enabled

```
#!/usr/bin/env bash

{
    l_output="" l_output2=""
    a_parlist=("net.ipv4.conf.all.accept_redirects=0" "net.ipv4.conf.default.accept_redirects=0"
"net.ipv6.conf.all.accept_redirects=0" "net.ipv6.conf.default.accept_redirects=0")
    l_ufwscf="$([ -f /etc/default/ufw ] && awk -F= '/^\s*IPT_SYSCTL=/{print $2}'
/etc/default/ufw)"
    kernel_parameter_chk()
    {
        l_krp="$(sysctl "$l_kpname" | awk -F= '{print $2}' | xargs)" # Check running configuration
        if [ "$l_krp" = "$l_kpvalue" ]; then
            l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_krp\" in the running
configuration"
        else
            l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_krp\" in the running
configuration and should have a value of: \"$l_kpvalue\""
        fi
        unset A_out; declare -A A_out # Check durable setting (files)
        while read -r l_out; do
            if [ -n "$l_out" ]; then
                if [[ $l_out =~ ^\s*# ]]; then
                    l_file="$l_out//# /"
                else
                    l_kpar="$(awk -F= '{print $1}' <<< "$l_out" | xargs)"
                    [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]="$l_file")
                fi
            fi
        done < <(/usr/lib/systemd/systemd-sysctl --cat-config | grep -Po
'^\h*([^\#\n\r]+|#\h*\n\r|)\.conf\b)')
        if [ -n "$l_ufwscf" ]; then # Account for systems with UFW (Not covered by systemd-sysctl -
-cat-config)
            l_kpar="$(grep -Po "\h*$l_kpname\b" "$l_ufwscf" | xargs)"
            l_kpar="$l_kpar//.//."
            [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]="$l_ufwscf")
        fi
        if (( ${#A_out[@]} > 0 )); then # Assess output from files and generate output
            while IFS="" read -r l_fkpname l_fkpvalue; do
                l_fkpname="$l_fkpname// /"; l_fkpvalue="$l_fkpvalue// /"
                if [ "$l_fkpvalue" = "$l_kpvalue" ]; then
                    l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_fkpvalue\" in
\"$(printf '%s' "${A_out[@]}")\""
                else
                    l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_fkpvalue\" in
\"$(printf '%s' "${A_out[@]}")\" and should have a value of: \"$l_kpvalue\""
                fi
            done < <(grep -Po -- "\h*$l_kpname\h*=\h*\H+" "${A_out[@]})
        else
            l_output2="$l_output2\n - \"$l_kpname\" is not set in an included file\n    ** Note:
\"$l_kpname\" May be set in a file that's ignored by load procedure **\n"
        fi
    }
    while IFS="" read -r l_kpname l_kpvalue; do # Assess and check parameters
        l_kpname="$l_kpname// /"; l_kpvalue="$l_kpvalue// /"
        if ! grep -Pqs '\h*0\b' /sys/module/ipv6/parameters/disable && grep -q '^net.ipv6.' <<<
"$l_kpname"; then
            l_output="$l_output\n - IPv6 is disabled on the system, \"$l_kpname\" is not applicable"
        else
            kernel_parameter_chk
        fi
    done < <(printf '%s\n' "${a_parlist[@]}")
    if [ -z "$l_output2" ]; then # Provide output from checks
        echo -e "\n- Audit Result:\n    ** PASS **\n$l_output\n"
    else
        echo -e "\n- Audit Result:\n    ** FAIL **\n - Reason(s) for audit failure:\n$l_output2\n"
        [ -n "$l_output" ] && echo -e "\n- Correctly set:\n$l_output\n"
    fi
}

```

Remediation:

Set the following parameters in `/etc/sysctl.conf` or a file in `/etc/sysctl.d/` ending in `.conf`:

- `net.ipv4.conf.all.accept_redirects = 0`
- `net.ipv4.conf.default.accept_redirects = 0`

Example:

```
# printf "  
net.ipv4.conf.all.accept_redirects = 0  
net.ipv4.conf.default.accept_redirects = 0  
" >> /etc/sysctl.d/60-netipv4_sysctl.conf
```

Run the following command to set the active kernel parameters:

```
# {  
    sysctl -w net.ipv4.conf.all.accept_redirects=0  
    sysctl -w net.ipv4.conf.default.accept_redirects=0  
    sysctl -w net.ipv4.route.flush=1  
}
```

-IF- IPv6 is enabled on the system:

Set the following parameters in `/etc/sysctl.conf` or a file in `/etc/sysctl.d/` ending in `.conf`:

- `net.ipv6.conf.all.accept_redirects = 0`
- `net.ipv6.conf.default.accept_redirects = 0`

Example:

```
# printf "  
net.ipv6.conf.all.accept_redirects = 0  
net.ipv6.conf.default.accept_redirects = 0  
" >> /etc/sysctl.d/60-netipv6_sysctl.conf
```

Run the following command to set the active kernel parameters:

```
# {  
    sysctl -w net.ipv6.conf.all.accept_redirects=0  
    sysctl -w net.ipv6.conf.default.accept_redirects=0  
    sysctl -w net.ipv6.route.flush=1  
}
```

Note: If these settings appear in a conically later file, or later in the same file, these settings will be overwritten

Default Value:

`net.ipv4.conf.all.accept_redirects = 1`

`net.ipv4.conf.default.accept_redirects = 1`

`net.ipv6.conf.all.accept_redirects = 1`

`net.ipv6.conf.default.accept_redirects = 1`

References:

1. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 <u>Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●
v7	9.2 <u>Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.		●	●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1557, T1557.000	TA0006, TA0009	M1030, M1042

3.3.3 *Ensure secure ICMP redirects are not accepted* (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Secure ICMP redirects are the same as ICMP redirects, except they come from gateways listed on the default gateway list. It is assumed that these gateways are known to your system, and that they are likely to be secure.

Rationale:

It is still possible for even known gateways to be compromised. Setting `net.ipv4.conf.all.secure_redirects` and `net.ipv4.conf.default.secure_redirects` to 0 protects the system from routing table updates by possibly compromised known gateways.

Audit:

Run the following script to verify the following kernel parameters are set in the running configuration and correctly loaded from a kernel parameter configuration file:

- `net.ipv4.conf.all.secure_redirects` is set to 0
- `net.ipv4.conf.default.secure_redirects` is set to 0

Note: kernel parameters are loaded by file and parameter order precedence. The following script observes this precedence as part of the auditing procedure. The parameters being checked may be set correctly in a file. If that file is superseded, the parameter is overridden by an incorrect setting later in that file, or in a conically later file, that "correct" setting will be ignored both by the script and by the system during a normal kernel parameter load sequence.

```
#!/usr/bin/env bash

{
    l_output="" l_output2=""
    a_parlist=("net.ipv4.conf.all.secure_redirects=0" "net.ipv4.conf.default.secure_redirects=0")
    l_ufwscf="$([ -f /etc/default/ufw ] && awk -F= '/^\s*IPT_SYSCTL=/ {print $2}'
/etc/default/ufw)"
    kernel_parameter_chk()
    {
        l_krp="$(sysctl "$l_kpname" | awk -F= '{print $2}' | xargs)" # Check running configuration
        if [ "$l_krp" = "$l_kpvalue" ]; then
            l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_krp\" in the running
configuration"
        else
            l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_krp\" in the running
configuration and should have a value of: \"$l_kpvalue\""
        fi
        unset A_out; declare -A A_out # Check durable setting (files)
        while read -r l_out; do
            if [ -n "$l_out" ]; then
                if [[ $l_out =~ ^\s*# ]]; then
                    l_file="$l_out// # /)"
                else
                    l_kpar="$(awk -F= '{print $1}' <<< "$l_out" | xargs)"
                    [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]=" $l_file")
                fi
            fi
            done < <(/usr/lib/systemd/systemd-sysctl --cat-config | grep -Po
'^\h*([^\#\n\r]+|#\h*\[/^\#\n\r\h]+\.\conf\b)')
            if [ -n "$l_ufwscf" ]; then # Account for systems with UFW (Not covered by systemd-sysctl -
-cat-config)
                l_kpar="$(grep -Po '^ \h*$l_kpname\b' "$l_ufwscf" | xargs)"
                l_kpar="$l_kpar//\./)"
                [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]=" $l_ufwscf")
            fi
            if (( ${#A_out[@]} > 0 )); then # Assess output from files and generate output
                while IFS="" read -r l_fkpname l_fkpvalue; do
                    l_fkpname="$l_fkpname// /)"; l_fkpvalue="$l_fkpvalue// /)"
                    if [ "$l_fkpvalue" = "$l_kpvalue" ]; then
                        l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_fkpvalue\" in
\"$(printf '%s' "${A_out[@]}")\""\n"
                    else
                        l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_fkpvalue\" in
\"$(printf '%s' "${A_out[@]}")\"" and should have a value of: \"$l_kpvalue\""\n"
                    fi
                    done < <(grep -Po -- "\h*$l_kpname\h*=\h*\H+" "${A_out[@]}")
                else
                    l_output2="$l_output2\n - \"$l_kpname\" is not set in an included file\n    ** Note:
\"$l_kpname\" May be set in a file that's ignored by load procedure **\n"
                fi
            }
            while IFS="" read -r l_kpname l_kpvalue; do # Assess and check parameters
                l_kpname="$l_kpname// /)"; l_kpvalue="$l_kpvalue// /)"
                if ! grep -Pqs '\h*0\b' /sys/module/ipv6/parameters/disable && grep -q '^net.ipv6.' <<<
"$l_kpname"; then
                    l_output="$l_output\n - IPv6 is disabled on the system, \"$l_kpname\" is not applicable"
                else
                    kernel_parameter_chk
                fi
                done < <(printf '%s\n' "${a_parlist[@]}")
                if [ -z "$l_output2" ]; then # Provide output from checks
                    echo -e "\n- Audit Result:\n    ** PASS **\n$l_output\n"
                else
                    echo -e "\n- Audit Result:\n    ** FAIL **\n - Reason(s) for audit failure:\n$l_output2\n"
                    [ -n "$l_output" ] && echo -e "\n- Correctly set:\n$l_output\n"
                fi
            fi
        }
    }
}
```


Remediation:

Set the following parameters in `/etc/sysctl.conf` or a file in `/etc/sysctl.d/` ending in `.conf`:

- `net.ipv4.conf.all.secure_redirects = 0`
- `net.ipv4.conf.default.secure_redirects = 0`

Example:

```
# printf "  
net.ipv4.conf.all.secure_redirects = 0  
net.ipv4.conf.default.secure_redirects = 0  
" >> /etc/sysctl.d/60-netipv4_sysctl.conf
```

Run the following commands to set the active kernel parameters:

```
# {  
    sysctl -w net.ipv4.conf.all.secure_redirects=0  
    sysctl -w net.ipv4.conf.default.secure_redirects=0  
    sysctl -w net.ipv4.route.flush=1  
}
```

Note: If these settings appear in a conically later file, or later in the same file, these settings will be overwritten

Default Value:

`net.ipv4.conf.all.secure_redirects = 1`

`net.ipv4.conf.default.secure_redirects = 1`

References:

1. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.		●	●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1557, T1557.000	TA0006, TA0009	M1030, M1042

3.3.4 *Ensure suspicious packets are logged (Automated)*

Profile Applicability:

- Level 1 - Server

Description:

When enabled, this feature logs packets with un-routable source addresses to the kernel log.

Rationale:

Setting `net.ipv4.conf.all.log_martians` and `net.ipv4.conf.default.log_martians` to `1` enables this feature. Logging these packets allows an administrator to investigate the possibility that an attacker is sending spoofed packets to their system.

Audit:

Run the following script to verify the following kernel parameters are set in the running configuration and correctly loaded from a kernel parameter configuration file:

- `net.ipv4.conf.all.log_martians` is set to 1
- `net.ipv4.conf.default.log_martians` is set to 1

Note: kernel parameters are loaded by file and parameter order precedence. The following script observes this precedence as part of the auditing procedure. The parameters being checked may be set correctly in a file. If that file is superseded, the parameter is overridden by an incorrect setting later in that file, or in a conically later file, that "correct" setting will be ignored both by the script and by the system during a normal kernel parameter load sequence.

```

#!/usr/bin/env bash

{
    l_output="" l_output2=""
    a_parlist=("net.ipv4.conf.all.log_martians=1" "net.ipv4.conf.default.log_martians=1")
    l_ufwscf="$([ -f /etc/default/ufw ] && awk -F= '/^\s*IPT_SYSCTL=/ {print $2}'
/etc/default/ufw)"
    kernel_parameter_chk()
    {
        l_krp="$(sysctl "$l_kpname" | awk -F= '{print $2}' | xargs)" # Check running configuration
        if [ "$l_krp" = "$l_kpvalue" ]; then
            l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_krp\" in the running
configuration"
        else
            l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_krp\" in the running
configuration and should have a value of: \"$l_kpvalue\""
        fi
        unset A_out; declare -A A_out # Check durable setting (files)
        while read -r l_out; do
            if [ -n "$l_out" ]; then
                if [[ $l_out =~ ^\s*# ]]; then
                    l_file="$l_out// # /)"
                else
                    l_kpar="$(awk -F= '{print $1}' <<< "$l_out" | xargs)"
                    [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]=" $l_file")
                fi
            fi
            done < <(/usr/lib/systemd/systemd-sysctl --cat-config | grep -Po
'^\h*([^\#\n\r]+|#\h*\[/^\#\n\r\h]+\.\conf\b)')
            if [ -n "$l_ufwscf" ]; then # Account for systems with UFW (Not covered by systemd-sysctl -
-cat-config)
                l_kpar="$(grep -Po '^ \h*$l_kpname\b' "$l_ufwscf" | xargs)"
                l_kpar="$l_kpar//\./)"
                [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]=" $l_ufwscf")
            fi
            if (( ${#A_out[@]} > 0 )); then # Assess output from files and generate output
                while IFS="" read -r l_fkpname l_fkpvalue; do
                    l_fkpname="$l_fkpname// /)"; l_fkpvalue="$l_fkpvalue// /)"
                    if [ "$l_fkpvalue" = "$l_kpvalue" ]; then
                        l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_fkpvalue\" in
\"$(printf '%s' "${A_out[@]}")\""\n"
                    else
                        l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_fkpvalue\" in
\"$(printf '%s' "${A_out[@]}")\"" and should have a value of: \"$l_kpvalue\""\n"
                    fi
                    done < <(grep -Po -- "\^ \h*$l_kpname\h*=\h*\H+" "${A_out[@]}")
                else
                    l_output2="$l_output2\n - \"$l_kpname\" is not set in an included file\n    ** Note:
\"$l_kpname\" May be set in a file that's ignored by load procedure **\n"
                fi
            }
            while IFS="" read -r l_kpname l_kpvalue; do # Assess and check parameters
                l_kpname="$l_kpname// /)"; l_kpvalue="$l_kpvalue// /)"
                if ! grep -Pqs '\^ \h*0\b' /sys/module/ipv6/parameters/disable && grep -q '\^net.ipv6.' <<<
"$l_kpname"; then
                    l_output="$l_output\n - IPv6 is disabled on the system, \"$l_kpname\" is not applicable"
                else
                    kernel_parameter_chk
                fi
                done < <(printf '%s\n' "${a_parlist[@]}")
                if [ -z "$l_output2" ]; then # Provide output from checks
                    echo -e "\n- Audit Result:\n    ** PASS **\n $l_output\n"
                else
                    echo -e "\n- Audit Result:\n    ** FAIL **\n - Reason(s) for audit failure:\n $l_output2\n"
                    [ -n "$l_output" ] && echo -e "\n- Correctly set:\n $l_output\n"
                fi
            fi
        }
    }
}

```

Remediation:

Set the following parameters in `/etc/sysctl.conf` or a file in `/etc/sysctl.d/` ending in `.conf`:

- `net.ipv4.conf.all.log_martians = 1`
- `net.ipv4.conf.default.log_martians = 1`

Example:

```
# printf "  
net.ipv4.conf.all.log_martians = 1  
net.ipv4.conf.default.log_martians = 1  
" >> /etc/sysctl.d/60-netipv4_sysctl.conf
```

Run the following command to set the active kernel parameters:

```
# {  
    sysctl -w net.ipv4.conf.all.log_martians=1  
    sysctl -w net.ipv4.conf.default.log_martians=1  
    sysctl -w net.ipv4.route.flush=1  
}
```

Note: If these settings appear in a conically later file, or later in the same file, these settings will be overwritten

Default Value:

`net.ipv4.conf.all.log_martians = 0`

`net.ipv4.conf.default.log_martians = 0`

References:

1. NIST SP 800-53 Rev. 5: AU-3

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 <u>Collect Detailed Audit Logs</u> Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			
v7	6.3 <u>Enable Detailed Logging</u> Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0005	

3.3.5 Ensure broadcast ICMP requests are ignored (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Setting `net.ipv4.icmp_echo_ignore_broadcasts` to 1 will cause the system to ignore all ICMP echo and timestamp requests to broadcast and multicast addresses.

Rationale:

Accepting ICMP echo and timestamp requests with broadcast or multicast destinations for your network could be used to trick your host into starting (or participating) in a Smurf attack. A Smurf attack relies on an attacker sending large amounts of ICMP broadcast messages with a spoofed source address. All hosts receiving this message and responding would send echo-reply messages back to the spoofed address, which is probably not routable. If many hosts respond to the packets, the amount of traffic on the network could be significantly multiplied.

Audit:

Run the following script to verify the following kernel parameter is set in the running configuration and correctly loaded from a kernel parameter configuration file:

- `net.ipv4.icmp_echo_ignore_broadcasts` is set to 1

Note: kernel parameters are loaded by file and parameter order precedence. The following script observes this precedence as part of the auditing procedure. The parameters being checked may be set correctly in a file. If that file is superseded, the parameter is overridden by an incorrect setting later in that file, or in a conically later file, that "correct" setting will be ignored both by the script and by the system during a normal kernel parameter load sequence.

```

#!/usr/bin/env bash

{
    l_output="" l_output2=""
    a_parlist=("net.ipv4.icmp_echo_ignore_broadcasts=1")
    l_ufwscf="$( [ -f /etc/default/ufw ] && awk -F= '/^\s*IPT_SYSCTL=/ {print $2}'
/etc/default/ufw )"
    kernel_parameter_chk()
    {
        l_krp="$(sysctl "$l_kpname" | awk -F= '{print $2}' | xargs)" # Check running configuration
        if [ "$l_krp" = "$l_kpvalue" ]; then
            l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_krp\" in the running
configuration"
        else
            l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_krp\" in the running
configuration and should have a value of: \"$l_kpvalue\""
        fi
        unset A_out; declare -A A_out # Check durable setting (files)
        while read -r l_out; do
            if [ -n "$l_out" ]; then
                if [[ $l_out =~ ^\s*# ]]; then
                    l_file="$l_out// # /)"
                else
                    l_kpar="$(awk -F= '{print $1}' <<< "$l_out" | xargs)"
                    [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]=" $l_file")
                fi
            fi
            done < <(/usr/lib/systemd/systemd-sysctl --cat-config | grep -Po
'^\h*([^\#\n\r]+|#\h*\[/^\#\n\r\h]+\.\conf\b)')
            if [ -n "$l_ufwscf" ]; then # Account for systems with UFW (Not covered by systemd-sysctl -
-cat-config)
                l_kpar="$(grep -Po '^ \h*$l_kpname\b' "$l_ufwscf" | xargs)"
                l_kpar="$l_kpar//\./)"
                [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]=" $l_ufwscf")
            fi
            if (( ${#A_out[@]} > 0 )); then # Assess output from files and generate output
                while IFS="" read -r l_fkpname l_fkpvalue; do
                    l_fkpname="$l_fkpname// /)"; l_fkpvalue="$l_fkpvalue// /)"
                    if [ "$l_fkpvalue" = "$l_kpvalue" ]; then
                        l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_fkpvalue\" in
\"$(printf '%s' "${A_out[@]}")\""\n"
                    else
                        l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_fkpvalue\" in
\"$(printf '%s' "${A_out[@]}")\"" and should have a value of: \"$l_kpvalue\""\n"
                    fi
                    done < <(grep -Po -- "\^ \h*$l_kpname\h*=\h*\H+" "${A_out[@]}")
                else
                    l_output2="$l_output2\n - \"$l_kpname\" is not set in an included file\n    ** Note:
\"$l_kpname\" May be set in a file that's ignored by load procedure **\n"
                fi
            }
            while IFS="" read -r l_kpname l_kpvalue; do # Assess and check parameters
                l_kpname="$l_kpname// /)"; l_kpvalue="$l_kpvalue// /)"
                if ! grep -Pqs '\^ \h*0\b' /sys/module/ipv6/parameters/disable && grep -q '\^net.ipv6.' <<<
"$l_kpname"; then
                    l_output="$l_output\n - IPv6 is disabled on the system, \"$l_kpname\" is not applicable"
                else
                    kernel_parameter_chk
                fi
                done < <(printf '%s\n' "${a_parlist[@]}")
                if [ -z "$l_output2" ]; then # Provide output from checks
                    echo -e "\n- Audit Result:\n    ** PASS **\n $l_output\n"
                else
                    echo -e "\n- Audit Result:\n    ** FAIL **\n - Reason(s) for audit failure:\n $l_output2\n"
                    [ -n "$l_output" ] && echo -e "\n- Correctly set:\n $l_output\n"
                fi
            fi
        }
    }
}

```


Remediation:

Set the following parameter in `/etc/sysctl.conf` or a file in `/etc/sysctl.d/` ending in `.conf`:

- `net.ipv4.icmp_echo_ignore_broadcasts = 1`

Example:

```
# printf "  
net.ipv4.icmp_echo_ignore_broadcasts = 1  
" >> /etc/sysctl.d/60-netipv4_sysctl.conf
```

Run the following command to set the active kernel parameters:

```
# {  
    sysctl -w net.ipv4.icmp_echo_ignore_broadcasts=1  
    sysctl -w net.ipv4.route.flush=1  
}
```

Note: If these settings appear in a conically later file, or later in the same file, these settings will be overwritten

Default Value:

`net.ipv4.conf.default.log_martians = 0`

References:

- NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.		●	●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1498, T1498.001	TA0040	M1037

3.3.6 *Ensure bogus ICMP responses are ignored (Automated)*

Profile Applicability:

- Level 1 - Server

Description:

Setting `net.ipv4.icmp_ignore_bogus_error_responses` to 1 prevents the kernel from logging bogus responses (RFC-1122 non-compliant) from broadcast reframes, keeping file systems from filling up with useless log messages.

Rationale:

Some routers (and some attackers) will send responses that violate RFC-1122 and attempt to fill up a log file system with many useless error messages.

Audit:

Run the following script to verify the following kernel parameter is set in the running configuration and correctly loaded from a kernel parameter configuration file:

- `net.ipv4.icmp_ignore_bogus_error_responses` is set to 1

Note: kernel parameters are loaded by file and parameter order precedence. The following script observes this precedence as part of the auditing procedure. The parameters being checked may be set correctly in a file. If that file is superseded, the parameter is overridden by an incorrect setting later in that file, or in a conically later file, that "correct" setting will be ignored both by the script and by the system during a normal kernel parameter load sequence.

```

#!/usr/bin/env bash

{
    l_output="" l_output2=""
    a_parlist=("net.ipv4.icmp_ignore_bogus_error_responses=1")
    l_ufwscf="$([ -f /etc/default/ufw ] && awk -F= '/^s*IPT_SYSCTL=/ {print $2}'
/etc/default/ufw)"
    kernel_parameter_chk()
    {
        l_krp="$(sysctl "$l_kpname" | awk -F= '{print $2}' | xargs)" # Check running configuration
        if [ "$l_krp" = "$l_kpvalue" ]; then
            l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_krp\" in the running
configuration"
        else
            l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_krp\" in the running
configuration and should have a value of: \"$l_kpvalue\""
        fi
        unset A_out; declare -A A_out # Check durable setting (files)
        while read -r l_out; do
            if [ -n "$l_out" ]; then
                if [[ $l_out =~ ^s*# ]]; then
                    l_file="$l_out//# /)"
                else
                    l_kpar="$(awk -F= '{print $1}' <<< "$l_out" | xargs)"
                    [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]=" $l_file")
                fi
            fi
            done < <(/usr/lib/systemd/systemd-sysctl --cat-config | grep -Po
'^\h*([^\#\n\r]+|#\h*\[/^\#\n\r\h]+\.\conf\b)')
            if [ -n "$l_ufwscf" ]; then # Account for systems with UFW (Not covered by systemd-sysctl -
-cat-config)
                l_kpar="$(grep -Po '^h*$l_kpname\b' "$l_ufwscf" | xargs)"
                l_kpar="$l_kpar//\./)"
                [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]=" $l_ufwscf")
            fi
            if (( ${#A_out[@]} > 0 )); then # Assess output from files and generate output
                while IFS="" read -r l_fkpname l_fkpvalue; do
                    l_fkpname="$l_fkpname// /)"; l_fkpvalue="$l_fkpvalue// /)"
                    if [ "$l_fkpvalue" = "$l_kpvalue" ]; then
                        l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_fkpvalue\" in
\"$(printf '%s' "${A_out[@]}")\""\n"
                    else
                        l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_fkpvalue\" in
\"$(printf '%s' "${A_out[@]}")\" and should have a value of: \"$l_kpvalue\""\n"
                    fi
                    done < <(grep -Po -- "\^h*$l_kpname\h*=\h*\H+" "${A_out[@]}")
                else
                    l_output2="$l_output2\n - \"$l_kpname\" is not set in an included file\n    ** Note:
\"$l_kpname\" May be set in a file that's ignored by load procedure **\n"
                fi
            }
            while IFS="" read -r l_kpname l_kpvalue; do # Assess and check parameters
                l_kpname="$l_kpname// /)"; l_kpvalue="$l_kpvalue// /)"
                if ! grep -Pqs '\^h*0\b' /sys/module/ipv6/parameters/disable && grep -q '\^net.ipv6.' <<<
"$l_kpname"; then
                    l_output="$l_output\n - IPv6 is disabled on the system, \"$l_kpname\" is not applicable"
                else
                    kernel_parameter_chk
                fi
                done < <(printf '%s\n' "${a_parlist[@]}")
                if [ -z "$l_output2" ]; then # Provide output from checks
                    echo -e "\n- Audit Result:\n    ** PASS **\n$l_output\n"
                else
                    echo -e "\n- Audit Result:\n    ** FAIL **\n - Reason(s) for audit failure:\n$l_output2\n"
                    [ -n "$l_output" ] && echo -e "\n- Correctly set:\n$l_output\n"
                fi
            fi
        }
    }
}

```

Remediation:

Set the following parameter in `/etc/sysctl.conf` or a file in `/etc/sysctl.d/` ending in `.conf`:

- `net.ipv4.icmp_ignore_bogus_error_responses = 1`

Example:

```
# printf "  
net.ipv4.icmp_ignore_bogus_error_responses = 1  
" >> /etc/sysctl.d/60-netipv4_sysctl.conf
```

Run the following command to set the active kernel parameters:

```
# {  
    sysctl -w net.ipv4.icmp_ignore_bogus_error_responses=1  
    sysctl -w net.ipv4.route.flush=1  
}
```

Note: If these settings appear in a conically later file, or later in the same file, these settings will be overwritten

Default Value:

`net.ipv4.icmp_ignore_bogus_error_responses = 1`

References:

- NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 <u>Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●
v7	9.2 <u>Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.		●	●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0040	M1053

3.3.7 Ensure Reverse Path Filtering is enabled (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Setting `net.ipv4.conf.all.rp_filter` and `net.ipv4.conf.default.rp_filter` to 1 forces the Linux kernel to utilize reverse path filtering on a received packet to determine if the packet was valid. Essentially, with reverse path filtering, if the return packet does not go out the same interface that the corresponding source packet came from, the packet is dropped (and logged if `log_martians` is set).

Rationale:

Setting `net.ipv4.conf.all.rp_filter` and `net.ipv4.conf.default.rp_filter` to 1 is a good way to deter attackers from sending your system bogus packets that cannot be responded to. One instance where this feature breaks down is if asymmetrical routing is employed. This would occur when using dynamic routing protocols (bgp, ospf, etc) on your system. If you are using asymmetrical routing on your system, you will not be able to enable this feature without breaking the routing.

Impact:

If you are using asymmetrical routing on your system, you will not be able to enable this feature without breaking the routing.

Audit:

Run the following script to verify the following kernel parameters are set in the running configuration and correctly loaded from a kernel parameter configuration file:

- `net.ipv4.conf.all.rp_filter` is set to 1
- `net.ipv4.conf.default.rp_filter` is set to 1

Note: kernel parameters are loaded by file and parameter order precedence. The following script observes this precedence as part of the auditing procedure. The parameters being checked may be set correctly in a file. If that file is superseded, the parameter is overridden by an incorrect setting later in that file, or in a conically later file, that "correct" setting will be ignored both by the script and by the system during a normal kernel parameter load sequence.

```

#!/usr/bin/env bash

{
    l_output="" l_output2=""
    a_parlist=("net.ipv4.conf.all.rp_filter=1" "net.ipv4.conf.default.rp_filter=1")
    l_ufwscf="$([ -f /etc/default/ufw ] && awk -F= '/^\s*IPT_SYSCTL=/ {print $2}'
/etc/default/ufw)"
    kernel_parameter_chk()
    {
        l_krp="$(sysctl "$l_kpname" | awk -F= '{print $2}' | xargs)" # Check running configuration
        if [ "$l_krp" = "$l_kpvalue" ]; then
            l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_krp\" in the running
configuration"
        else
            l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_krp\" in the running
configuration and should have a value of: \"$l_kpvalue\""
        fi
        unset A_out; declare -A A_out # Check durable setting (files)
        while read -r l_out; do
            if [ -n "$l_out" ]; then
                if [[ $l_out =~ ^\s*# ]]; then
                    l_file="$l_out//# /)"
                else
                    l_kpar="$(awk -F= '{print $1}' <<< "$l_out" | xargs)"
                    [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]=" $l_file")
                fi
            fi
            done < <(/usr/lib/systemd/systemd-sysctl --cat-config | grep -Po
'^\h*([^\#\n\r]+|#\h*\[/^\#\n\r\h]+\.\conf\b)')
            if [ -n "$l_ufwscf" ]; then # Account for systems with UFW (Not covered by systemd-sysctl -
-cat-config)
                l_kpar="$(grep -Po '^ \h*$l_kpname\b' "$l_ufwscf" | xargs)"
                l_kpar="$l_kpar//\./)"
                [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]=" $l_ufwscf")
            fi
            if (( ${#A_out[@]} > 0 )); then # Assess output from files and generate output
                while IFS="" read -r l_fkpname l_fkpvalue; do
                    l_fkpname="$l_fkpname// /)"; l_fkpvalue="$l_fkpvalue// /)"
                    if [ "$l_fkpvalue" = "$l_kpvalue" ]; then
                        l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_fkpvalue\" in
\"$(printf '%s' "${A_out[@]}")\""\n"
                    else
                        l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_fkpvalue\" in
\"$(printf '%s' "${A_out[@]}")\"" and should have a value of: \"$l_kpvalue\""\n"
                    fi
                    done < <(grep -Po -- "\^ \h*$l_kpname\h*=\h*\H+" "${A_out[@]}")
                else
                    l_output2="$l_output2\n - \"$l_kpname\" is not set in an included file\n    ** Note:
\"$l_kpname\" May be set in a file that's ignored by load procedure **\n"
                fi
            }
            while IFS="" read -r l_kpname l_kpvalue; do # Assess and check parameters
                l_kpname="$l_kpname// /)"; l_kpvalue="$l_kpvalue// /)"
                if ! grep -Pqs '\^ \h*0\b' /sys/module/ipv6/parameters/disable && grep -q '\^net.ipv6.' <<<
"$l_kpname"; then
                    l_output="$l_output\n - IPv6 is disabled on the system, \"$l_kpname\" is not applicable"
                else
                    kernel_parameter_chk
                fi
                done < <(printf '%s\n' "${a_parlist[@]}")
                if [ -z "$l_output2" ]; then # Provide output from checks
                    echo -e "\n- Audit Result:\n    ** PASS **\n$l_output\n"
                else
                    echo -e "\n- Audit Result:\n    ** FAIL **\n - Reason(s) for audit failure:\n$l_output2\n"
                    [ -n "$l_output" ] && echo -e "\n- Correctly set:\n$l_output\n"
                fi
            fi
        }
    }
}

```


Remediation:

Set the following parameters in `/etc/sysctl.conf` or a file in `/etc/sysctl.d/` ending in `.conf`:

- `net.ipv4.conf.all.rp_filter = 1`
- `net.ipv4.conf.default.rp_filter = 1`

Example:

```
# printf "  
net.ipv4.conf.all.rp_filter = 1  
net.ipv4.conf.default.rp_filter = 1  
" >> /etc/sysctl.d/60-netipv4_sysctl.conf
```

Run the following commands to set the active kernel parameters:

```
# {  
    sysctl -w net.ipv4.conf.all.rp_filter=1  
    sysctl -w net.ipv4.conf.default.rp_filter=1  
    sysctl -w net.ipv4.route.flush=1  
}
```

Note: If these settings appear in a conically later file, or later in the same file, these settings will be overwritten

Default Value:

`net.ipv4.conf.all.rp_filter = 2`

`net.ipv4.conf.default.rp_filter = 1`

References:

1. NIST SP 800-53 Rev. 5: CM-1,CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.		●	●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1498, T1498.001	TA0006, TA0040	M1030, M1042

3.3.8 Ensure TCP SYN Cookies is enabled (Automated)

Profile Applicability:

- Level 1 - Server

Description:

When `tcp_syncookies` is set, the kernel will handle TCP SYN packets normally until the half-open connection queue is full, at which time, the SYN cookie functionality kicks in. SYN cookies work by not using the SYN queue at all. Instead, the kernel simply replies to the SYN with a SYN|ACK, but will include a specially crafted TCP sequence number that encodes the source and destination IP address and port number and the time the packet was sent. A legitimate connection would send the ACK packet of the three way handshake with the specially crafted sequence number. This allows the system to verify that it has received a valid response to a SYN cookie and allow the connection, even though there is no corresponding SYN in the queue.

Rationale:

Attackers use SYN flood attacks to perform a denial of service attacked on a system by sending many SYN packets without completing the three way handshake. This will quickly use up slots in the kernel's half-open connection queue and prevent legitimate connections from succeeding. Setting `net.ipv4.tcp_syncookies` to 1 enables SYN cookies, allowing the system to keep accepting valid connections, even if under a denial of service attack.

Audit:

Run the following script to verify the following kernel parameter is set in the running configuration and correctly loaded from a kernel parameter configuration file:

- `net.ipv4.tcp_syncookies` is set to 1

Note: kernel parameters are loaded by file and parameter order precedence. The following script observes this precedence as part of the auditing procedure. The parameters being checked may be set correctly in a file. If that file is superseded, the parameter is overridden by an incorrect setting later in that file, or in a conically later file, that "correct" setting will be ignored both by the script and by the system during a normal kernel parameter load sequence.

```

#!/usr/bin/env bash

{
    l_output="" l_output2=""
    a_parlist=("net.ipv4.tcp_syncookies=1")
    l_ufwscf="$( [ -f /etc/default/ufw ] && awk -F= '/^\s*IPT_SYSCTL=/ {print $2}'
/etc/default/ufw )"
    kernel_parameter_chk()
    {
        l_krp="$(sysctl "$l_kpname" | awk -F= '{print $2}' | xargs)" # Check running configuration
        if [ "$l_krp" = "$l_kpvalue" ]; then
            l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_krp\" in the running
configuration"
        else
            l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_krp\" in the running
configuration and should have a value of: \"$l_kpvalue\""
        fi
        unset A_out; declare -A A_out # Check durable setting (files)
        while read -r l_out; do
            if [ -n "$l_out" ]; then
                if [[ $l_out =~ ^\s*# ]]; then
                    l_file="$l_out// # /)"
                else
                    l_kpar="$(awk -F= '{print $1}' <<< "$l_out" | xargs)"
                    [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]=" $l_file")
                fi
            fi
            done < <(/usr/lib/systemd/systemd-sysctl --cat-config | grep -Po
'^\h*([^\#\n\r]+|#\h*\[/^\#\n\r\h]+\.\conf\b)')
            if [ -n "$l_ufwscf" ]; then # Account for systems with UFW (Not covered by systemd-sysctl -
-cat-config)
                l_kpar="$(grep -Po '^ \h*$l_kpname\b' "$l_ufwscf" | xargs)"
                l_kpar="$l_kpar//\./)"
                [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]=" $l_ufwscf")
            fi
            if (( ${#A_out[@]} > 0 )); then # Assess output from files and generate output
                while IFS="" read -r l_fkpname l_fkpvalue; do
                    l_fkpname="$l_fkpname// /)"; l_fkpvalue="$l_fkpvalue// /)"
                    if [ "$l_fkpvalue" = "$l_kpvalue" ]; then
                        l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_fkpvalue\" in
\"$(printf '%s' "${A_out[@]}")\""\n"
                    else
                        l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_fkpvalue\" in
\"$(printf '%s' "${A_out[@]}")\"" and should have a value of: \"$l_kpvalue\""\n"
                    fi
                    done < <(grep -Po -- "\^ \h*$l_kpname\h*=\h*\H+" "${A_out[@]}")
                else
                    l_output2="$l_output2\n - \"$l_kpname\" is not set in an included file\n    ** Note:
\"$l_kpname\" May be set in a file that's ignored by load procedure **\n"
                fi
            }
            while IFS="" read -r l_kpname l_kpvalue; do # Assess and check parameters
                l_kpname="$l_kpname// /)"; l_kpvalue="$l_kpvalue// /)"
                if ! grep -Pqs '\^ \h*0\b' /sys/module/ipv6/parameters/disable && grep -q '\^net.ipv6.' <<<
"$l_kpname"; then
                    l_output="$l_output\n - IPv6 is disabled on the system, \"$l_kpname\" is not applicable"
                else
                    kernel_parameter_chk
                fi
                done < <(printf '%s\n' "${a_parlist[@]}")
                if [ -z "$l_output2" ]; then # Provide output from checks
                    echo -e "\n- Audit Result:\n    ** PASS **\n $l_output\n"
                else
                    echo -e "\n- Audit Result:\n    ** FAIL **\n - Reason(s) for audit failure:\n $l_output2\n"
                    [ -n "$l_output" ] && echo -e "\n- Correctly set:\n $l_output\n"
                fi
            fi
        }
    }
}

```

Remediation:

Set the following parameter in `/etc/sysctl.conf` or a file in `/etc/sysctl.d/` ending in `.conf`:

- `net.ipv4.tcp_syncookies = 1`

Example:

```
# printf "  
net.ipv4.tcp_syncookies = 1  
" >> /etc/sysctl.d/60-netipv4_sysctl.conf
```

Run the following command to set the active kernel parameters:

```
# {  
    sysctl -w net.ipv4.tcp_syncookies=1  
    sysctl -w net.ipv4.route.flush=1  
}
```

Note: If these settings appear in a conically later file, or later in the same file, these settings will be overwritten

Default Value:

`net.ipv4.tcp_syncookies = 1`

References:

- NIST SP 800-53 Rev. 5: CM-1,CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.		●	●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1499, T1499.001	TA0040	M1037

3.3.9 *Ensure IPv6 router advertisements are not accepted (Automated)*

Profile Applicability:

- Level 1 - Server

Description:

This setting disables the system's ability to accept IPv6 router advertisements.

Rationale:

It is recommended that systems do not accept router advertisements as they could be tricked into routing traffic to compromised machines. Setting hard routes within the system (usually a single default route to a trusted router) protects the system from bad routes. Setting `net.ipv6.conf.all.accept_ra` and `net.ipv6.conf.default.accept_ra` to 0 disables the system's ability to accept IPv6 router advertisements.

Audit:

Run the following script to verify the following kernel parameters are set in the running configuration and correctly loaded from a kernel parameter configuration file:

- `net.ipv6.conf.all.accept_ra` is set to 0
- `net.ipv6.conf.default.accept_ra` is set to 0

Note:

- kernel parameters are loaded by file and parameter order precedence. The following script observes this precedence as part of the auditing procedure. The parameters being checked may be set correctly in a file. If that file is superseded, the parameter is overridden by an incorrect setting later in that file, or in a conically later file, that "correct" setting will be ignored both by the script and by the system during a normal kernel parameter load sequence.
- IPv6 kernel parameters only apply to systems where IPv6 is enabled

```

#!/usr/bin/env bash

{
    l_output="" l_output2=""
    a_parlist=("net.ipv6.conf.all.accept_ra=0" "net.ipv6.conf.default.accept_ra=0")
    l_ufwscf="$([ -f /etc/default/ufw ] && awk -F= '/^\s*IPT_SYSCTL=/ {print $2}'
/etc/default/ufw)"
    kernel_parameter_chk()
    {
        l_krp="$(sysctl "$l_kpname" | awk -F= '{print $2}' | xargs)" # Check running configuration
        if [ "$l_krp" = "$l_kpvalue" ]; then
            l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_krp\" in the running
configuration"
        else
            l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_krp\" in the running
configuration and should have a value of: \"$l_kpvalue\""
        fi
        unset A_out; declare -A A_out # Check durable setting (files)
        while read -r l_out; do
            if [ -n "$l_out" ]; then
                if [[ $l_out =~ ^\s*# ]]; then
                    l_file="$l_out//# /)"
                else
                    l_kpar="$(awk -F= '{print $1}' <<< "$l_out" | xargs)"
                    [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]=" $l_file")
                fi
            fi
            done < <(/usr/lib/systemd/systemd-sysctl --cat-config | grep -Po
'^\h*([^\#\n\r]+|#\h*\[/^\#\n\r\h]+\.\conf\b)')
            if [ -n "$l_ufwscf" ]; then # Account for systems with UFW (Not covered by systemd-sysctl -
-cat-config)
                l_kpar="$(grep -Po '^h*$l_kpname\b' "$l_ufwscf" | xargs)"
                l_kpar="$l_kpar//\./)"
                [ "$l_kpar" = "$l_kpname" ] && A_out+=(["$l_kpar"]=" $l_ufwscf")
            fi
            if (( ${#A_out[@]} > 0 )); then # Assess output from files and generate output
                while IFS="" read -r l_fkpname l_fkpvalue; do
                    l_fkpname="$l_fkpname// /)"; l_fkpvalue="$l_fkpvalue// /)"
                    if [ "$l_fkpvalue" = "$l_kpvalue" ]; then
                        l_output="$l_output\n - \"$l_kpname\" is correctly set to \"$l_fkpvalue\" in
\"$(printf '%s' "${A_out[@]}")\""\n"
                    else
                        l_output2="$l_output2\n - \"$l_kpname\" is incorrectly set to \"$l_fkpvalue\" in
\"$(printf '%s' "${A_out[@]}")\"" and should have a value of: \"$l_kpvalue\""\n"
                    fi
                    done < <(grep -Po -- "\^h*$l_kpname\h*=\h*\H+" "${A_out[@]}")
                else
                    l_output2="$l_output2\n - \"$l_kpname\" is not set in an included file\n    ** Note:
\"$l_kpname\" May be set in a file that's ignored by load procedure **\n"
                fi
            }
            while IFS="" read -r l_kpname l_kpvalue; do # Assess and check parameters
                l_kpname="$l_kpname// /)"; l_kpvalue="$l_kpvalue// /)"
                if ! grep -Pqs '\^h*0\b' /sys/module/ipv6/parameters/disable && grep -q '\^net.ipv6.' <<<
"$l_kpname"; then
                    l_output="$l_output\n - IPv6 is disabled on the system, \"$l_kpname\" is not applicable"
                else
                    kernel_parameter_chk
                fi
                done < <(printf '%s\n' "${a_parlist[@]}")
                if [ -z "$l_output2" ]; then # Provide output from checks
                    echo -e "\n- Audit Result:\n    ** PASS **\n$l_output\n"
                else
                    echo -e "\n- Audit Result:\n    ** FAIL **\n - Reason(s) for audit failure:\n$l_output2\n"
                    [ -n "$l_output" ] && echo -e "\n- Correctly set:\n$l_output\n"
                fi
            fi
        }
    }
}

```


Remediation:

-IF- IPv6 is enabled on the system:

Set the following parameters in `/etc/sysctl.conf` or a file in `/etc/sysctl.d/` ending in `.conf`:

- `net.ipv6.conf.all.accept_ra = 0`
- `net.ipv6.conf.default.accept_ra = 0`

Example:

```
# printf "  
net.ipv6.conf.all.accept_ra = 0  
net.ipv6.conf.default.accept_ra = 0  
" >> /etc/sysctl.d/60-netipv6_sysctl.conf
```

Run the following command to set the active kernel parameters:

```
# {  
  sysctl -w net.ipv6.conf.all.accept_ra=0  
  sysctl -w net.ipv6.conf.default.accept_ra=0  
  sysctl -w net.ipv6.route.flush=1  
}
```

Note: If these settings appear in a conically later file, or later in the same file, these settings will be overwritten

Default Value:

`net.ipv6.conf.all.accept_ra = 1`

`net.ipv6.conf.default.accept_ra = 1`

References:

1. NIST SP 800-53 Rev. 5: CM-1,CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.		●	●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1557, T1557.000	TA0006, TA0040	M1030, M1042

3.4 Configure Host Based Firewall

Amazon Elastic Compute Cloud security groups enable you to control traffic to your instance, including the kind of traffic that can reach your instance. For example, you can allow computers from only your home network to access your instance using SSH.

A Host Based Firewall, on a Linux system, is a set of rules. When a data packet moves into or out of a protected network space, its contents (in particular, information about its origin, target, and the protocol it plans to use) are tested against the firewall rules to see if it should be allowed through

To provide a Host Based Firewall, the Linux kernel includes support for nftables.

- nftables - A subsystem of the Linux kernel providing filtering and classification of network packets/datagrams/frames. nftables is supposed to replace certain parts of Netfilter, while keeping and reusing most of it. nftables utilizes the building blocks of the Netfilter infrastructure, such as the existing hooks into the networking stack, connection tracking system, userspace queueing component, and logging subsystem.

In order to configure firewall rules for nftables, a firewall utility needs to be installed. Guidance has been included for the following firewall utilities:

- FirewallD - Provides firewall features by acting as a front-end for the Linux kernel's netfilter framework via the nftables backend. Starting in v0.6.0, FirewallD added support for acting as a front-end for the Linux kernel's netfilter framework via the nftables userspace utility, acting as an alternative to the nft command line program. firewalld supports both IPv4 and IPv6 networks and can administer separate firewall zones with varying degrees of trust as defined in zone profiles.
- nftables - Includes the nft utility for configuration of the nftables subsystem of the Linux kernel

Note:

- **If security groups are configured and meet local site security policies, this section may be skipped**
- Only **one** method should be used to configure a firewall on the system. Use of more than one method could produce unexpected results.
- This section is intended only to ensure the resulting firewall rules are in place, not how they are configured.
- The `ipset` and `iptables-nft` packages have been deprecated in Fedora 34 based Linux distributions. This includes deprecation of nft-variants such as `iptables`, `ip6tables`, `arptables`, and `ebtables` utilities. If you are using any of these tools, for example, because you upgraded from an earlier version, we recommend migrating to the nft command line tool provided by the nftables package.

3.4.1 Configure a firewall utility

In order to configure firewall rules for Netfilter or nftables, a firewall utility needs to be installed. Guidance has been included for the following firewall utilities:

- FirewallD:
 - Provides firewall features by acting as a front-end for the Linux kernel's netfilter framework via the nftables userspace utility, acting as an alternative to the nft command line program. firewalld supports both IPv4 and IPv6 networks and can administer separate firewall zones with varying degrees of trust as defined in zone profiles.
 - Use the firewalld utility for simple firewall use cases. The utility is easy to use and covers the typical use cases for these scenarios.
- NfTables:
 - Includes the nft utility for configuration of the nftables subsystem of the Linux kernel
 - Use the nftables utility to set up complex and performance critical firewalls, such as for a whole network.

Note:

- **If security groups are configured and meet local site security policies, this section may be skipped**
- Only **one** method should be used to configure a firewall on the system. Use of more than one method could produce unexpected results.
- IPTables are deprecated in this release, and not covered in this Benchmark. The iptables utility on Fedora 34/CentOS 9 stream based Linux distributions uses the nf_tables kernel API instead of the legacy back end. The nf_tables API provides backward compatibility so that scripts that use iptables commands still work on Fedora 34 based Linux distributions. For new firewall scripts, it is recommended to use nftables.

3.4.1.1 Ensure *nftables* is installed (Automated)

Profile Applicability:

- Level 2 - Server

Description:

nftables provides a new in-kernel packet classification framework that is based on a network-specific Virtual Machine (VM) and a new *nft* userspace command line tool.

nftables reuses the existing Netfilter subsystems such as the existing hook infrastructure, the connection tracking system, NAT, userspace queuing and logging subsystem.

Rationale:

nftables is a subsystem of the Linux kernel that can protect against threats originating from within a corporate network to include malicious mobile code and poorly configured software on a host.

Impact:

Changing firewall settings while connected over the network can result in being locked out of the system.

Audit:

Run the following command to verify that *nftables* is installed:

```
# rpm -q nftables  
nftables-<version>
```

Remediation:

Run the following command to install *nftables*

```
# dnf install nftables
```

References:

1. NIST SP 800-53 Rev. 5: CM-7, CA-9

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.4 Implement and Manage a Firewall on Servers Implement and manage a firewall on servers, where supported. Example implementations include a virtual firewall, operating system firewall, or a third-party firewall agent.			
v7	9.4 Apply Host-based Firewalls or Port Filtering Apply host-based firewalls or port filtering tools on end systems, with a default-deny rule that drops all traffic except those services and ports that are explicitly allowed.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.004	TA0011	M1031, M1037

3.4.1.2 Ensure a single firewall configuration utility is in use (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Firewalld - Is a firewall service daemon that provides a dynamic customizable host-based firewall with a D-Bus interface. Being dynamic, it enables creating, changing, and deleting the rules without the necessity to restart the firewall daemon each time the rules are changed

NFTables - Includes the nft utility for configuration of the nftables subsystem of the Linux kernel

Note: firewalld with nftables backend does not support passing custom nftables rules to firewalld, using the `--direct` option.

Rationale:

In order to configure firewall rules for nftables, a firewall utility needs to be installed and active of the system. The use of more than one firewall utility may produce unexpected results.

Audit:

Run the following script to verify that a single firewall utility is in use on the system:

```
#!/usr/bin/env bash

{
    l_output="" l_output2="" l_fwd_status="" l_nft_status=""
    l_fwutil_status=""
    # Determine FirewallD utility Status
    rpm -q firewalld > /dev/null 2>&1 && l_fwd_status="$(systemctl is-enabled
firewalld.service):$(systemctl is-active firewalld.service)"
    # Determine NFTables utility Status
    rpm -q nftables > /dev/null 2>&1 && l_nft_status="$(systemctl is-enabled
nftables.service):$(systemctl is-active nftables.service)"
    l_fwutil_status="$l_fwd_status:$l_nft_status"
    case $l_fwutil_status in
        enabled:active:masked:inactive|enabled:active:disabled:inactive)
            l_output="\n - FirewallD utility is in use, enabled and active\n -
NFTables utility is correctly disabled or masked and inactive" ;;
        masked:inactive:enabled:active|disabled:inactive:enabled:active)
            l_output="\n - NFTables utility is in use, enabled and active\n -
FirewallD utility is correctly disabled or masked and inactive" ;;
        enabled:active:enabled:active)
            l_output2="\n - Both FirewallD and NFTables utilities are enabled
and active" ;;
        enabled*:enabled:*)
            l_output2="\n - Both FirewallD and NFTables utilities are enabled"
;;
        *:active*:active)
            l_output2="\n - Both FirewallD and NFTables utilities are enabled"
;;
        :enabled:active)
            l_output="\n - NFTables utility is in use, enabled, and active\n -
FirewallD package is not installed" ;;
        :)
            l_output2="\n - Neither FirewallD or NFTables is installed." ;;
        *:*:*)
            l_output2="\n - NFTables package is not installed on the system" ;;
        *)
            l_output2="\n - Unable to determine firewall state" ;;
    esac
    if [ -z "$l_output2" ]; then
        echo -e "\n- Audit Results:\n ** Pass **\n$l_output\n"
    else
        echo -e "\n- Audit Results:\n ** Fail **\n$l_output2\n"
    fi
}
```


Remediation:

Run the following script to ensure that a single firewall utility is in use on the system:

```
#!/usr/bin/env bash

{
  l_output="" l_output2="" l_fwd_status="" l_nft_status="" l_fwutil_status=""
  # Determine FirewallD utility Status
  rpm -q firewalld > /dev/null 2>&1 && l_fwd_status="$(systemctl is-enabled
firewalld.service):$(systemctl is-active firewalld.service)"
  # Determine Nftables utility Status
  rpm -q nftables > /dev/null 2>&1 && l_nft_status="$(systemctl is-enabled
nftables.service):$(systemctl is-active nftables.service)"
  l_fwutil_status="$l_fwd_status:$l_nft_status"
  case $l_fwutil_status in
    enabled:active:masked:inactive|enabled:active:disabled:inactive)
      echo -e "\n - FirewallD utility is in use, enabled and active\n - Nftables utility is
correctly disabled or masked and inactive\n - no remediation required" ;;
    masked:inactive:enabled:active|disabled:inactive:enabled:active)
      echo -e "\n - Nftables utility is in use, enabled and active\n - FirewallD utility is
correctly disabled or masked and inactive\n - no remediation required" ;;
    enabled:active:enabled:active)
      echo -e "\n - Both FirewallD and Nftables utilities are enabled and active\n - stopping
and masking Nftables utility"
      systemctl stop nftables && systemctl --now mask nftables ;;
    enabled*:enabled:*)
      echo -e "\n - Both FirewallD and Nftables utilities are enabled\n - remediating"
      if [ "$(awk -F: '{print $2}' <<< "$l_fwutil_status")" = "active" ] && [ "$(awk -F:
'{print $4}' <<< "$l_fwutil_status")" = "inactive" ]; then
        echo " - masking Nftables utility"
        systemctl stop nftables && systemctl --now mask nftables
      elif [ "$(awk -F: '{print $4}' <<< "$l_fwutil_status")" = "active" ] && [ "$(awk -F:
'{print $2}' <<< "$l_fwutil_status")" = "inactive" ]; then
        echo " - masking FirewallD utility"
        systemctl stop firewalld && systemctl --now mask firewalld
      fi ;;
    *:active*:active)
      echo -e "\n - Both FirewallD and Nftables utilities are active\n - remediating"
      if [ "$(awk -F: '{print $1}' <<< "$l_fwutil_status")" = "enabled" ] && [ "$(awk -F:
'{print $3}' <<< "$l_fwutil_status")" != "enabled" ]; then
        echo " - stopping and masking Nftables utility"
        systemctl stop nftables && systemctl --now mask nftables
      elif [ "$(awk -F: '{print $3}' <<< "$l_fwutil_status")" = "enabled" ] && [ "$(awk -F:
'{print $1}' <<< "$l_fwutil_status")" != "enabled" ]; then
        echo " - stopping and masking FirewallD utility"
        systemctl stop firewalld && systemctl --now mask firewalld
      fi ;;
    :enabled:active)
      echo -e "\n - Nftables utility is in use, enabled, and active\n - FirewallD package is
not installed\n - no remediation required" ;;
    :)
      echo -e "\n - Neither FirewallD or Nftables is installed.\n - remediating\n - installing
Nftables"
      dnf -q install nftables ;;
    *:*)
      echo -e "\n - Nftables package is not installed on the system\n - remediating\n -
installing Nftables"
      dnf -q install nftables ;;
    *)
      echo -e "\n - Unable to determine firewall state" ;;
  esac
}
```

References:

1. https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html-single/configuring_firewalls_and_packet_filters/index

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.4 <u>Implement and Manage a Firewall on Servers</u> Implement and manage a firewall on servers, where supported. Example implementations include a virtual firewall, operating system firewall, or a third-party firewall agent.			
v8	4.5 <u>Implement and Manage a Firewall on End-User Devices</u> Implement and manage a host-based firewall or port-filtering tool on end-user devices, with a default-deny rule that drops all traffic except those services and ports that are explicitly allowed.			
v7	9.4 <u>Apply Host-based Firewalls or Port Filtering</u> Apply host-based firewalls or port filtering tools on end systems, with a default-deny rule that drops all traffic except those services and ports that are explicitly allowed.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.004	TA0011	M1031, M1037

3.4.2 Configure firewall rules

Note: If security groups are configured and meet local site security policies, this section may be skipped

nftables is a subsystem of the Linux kernel providing filtering and classification of network packets/datagrams/frames and is the successor to iptables.

Important: Configuration of a live systems firewall directly over a remote connection will often result in being locked out. It is advised to have a known good firewall configuration set to run on boot and to configure an entire firewall structure in a script that is then run and tested before saving to boot.

Firewalld

The following will create a Firewalld Zone called `securezone` to implement the firewall rules of this section leveraging the `firewalld` utility included with the `firewalld` package. This example will open port 22(ssh) from anywhere. Opening service `SSH` should be updated in accordance with local site policy.

Sample securezone zone xml file

```
<?xml version="1.0" encoding="utf-8"?>
<zone target="DROP">
  <description>For use with CIS Linux Benchmark. You do not trust the other
  computers on networks to not harm your computer. Only selected incoming
  connections are accepted.</description>
  <service name="ssh"/>
  <service name="dhcpv6-client"/>
  <icmp-block name="destination-unreachable"/>
  <icmp-block name="packet-too-big"/>
  <icmp-block name="time-exceeded"/>
  <icmp-block name="parameter-problem"/>
  <icmp-block name="neighbour-advertisement"/>
  <icmp-block name="neighbour-solicitation"/>
  <icmp-block name="router-advertisement"/>
  <icmp-block name="router-solicitation"/>
  <rule family="ipv4">
    <source address="127.0.0.1"/>
    <destination address="127.0.0.1" invert="True"/>
    <drop/>
  </rule>
  <rule family="ipv6">
    <source address="::1"/>
    <destination address="::1" invert="True"/>
    <drop/>
  </rule>
  <icmp-block-inversion/>
</zone>
```

Note: To use this zone, save this as `/etc/firewalld/zones/securezone.xml` and run the following commands:

```
# firewall-cmd --reload
# firewall-cmd --permanent --zone=securezone --change-interface={NAME OF
NETWORK INTERFACE}
```

NFTables Utility:

The following will implement the firewall rules of this section leveraging the nftables utility included with the nftables package. This example will open ICMP, IGMP, and port 22(ssh) from anywhere. Opening the ports for ICMP, IGMP, and port 22(ssh) needs to be updated in accordance with local site policy. Allow port 22(ssh) should to be updated to only allow systems requiring ssh connectivity to connect, as per site policy.

Save the script below as `/etc/nftables/nftables.rules`

```

# This nftables.rules config should be saved as /etc/nftables/nftables.rules

# flush nftables ruleset
flush ruleset

# Load nftables ruleset
# nftables config with inet table named filter

table inet filter {
    chain input {
        type filter hook input priority 0; policy drop;

        # early drop invalid packets
        ct state invalid drop

        # allow loopback if not forged
        iif lo accept
        iif != lo ip daddr 127.0.0.1/8 drop
        iif != lo ip6 daddr ::1/128 drop

        # allow connections made by ourselves
        ip protocol tcp ct state established accept
        ip protocol udp ct state established accept

        # allow from anywhere
        ip protocol igmp accept
        tcp dport ssh accept

        # allow some icmp
        icmpv6 type { destination-unreachable, packet-too-big, time-exceeded,
parameter-problem, mld-listener-query, mld-listener-report, mld-listener-
done, nd-router-solicit, nd-router-advert, nd-neighbor-solicit, nd-neighbor-
advert, ind-neighbor-solicit, ind-neighbor-advert, mld2-listener-report }
accept
        icmp type { destination-unreachable, router-advertisement, router-
solicitation, time-exceeded, parameter-problem } accept
    }

    chain forward {
        # drop all forward
        type filter hook forward priority 0; policy drop;
    }

    chain output {
        # can omit this as its accept by default
        type filter hook output priority 0; policy accept;
    }
}

```

Run the following command to load the file into nftables

```
# nft -f /etc/nftables/nftables.rules
```

Note: All changes in the nftables subsections are temporary

To make these changes permanent:

Run the following command to create the nftables.rules file

```
nft list ruleset > /etc/nftables/nftables.rules
```

Add the following line to /etc/sysconfig/nftables.conf

```
include "/etc/nftables/nftables.rules"
```

3.4.2.1 Ensure firewalld default zone is set (Automated)

Profile Applicability:

- Level 2 - Server

Description:

A firewall zone defines the trust level for a connection, interface or source address binding. This is a one to many relation, which means that a connection, interface or source can only be part of one zone, but a zone can be used for many network connections, interfaces and sources.

- The default zone is the zone that is used for everything that is not explicitly bound/assigned to another zone.
- If no zone assigned to a connection, interface or source, only the default zone is used.
- The default zone is not always listed as being used for an interface or source as it will be used for it either way. This depends on the manager of the interfaces.

Connections handled by NetworkManager are listed as NetworkManager requests to add the zone binding for the interface used by the connection. Also interfaces under control of the network service are listed also because the service requests it.

Note:

- A firewalld zone configuration file contains the information for a zone.
 - These are the zone description, services, ports, protocols, icmp-blocks, masquerade, forward-ports and rich language rules in an XML file format.
 - The file name has to be `zone_name.xml` where length of `zone_name` is currently limited to 17 chars.
- NetworkManager binds interfaces to zones automatically

Rationale:

Because the default zone is the zone that is used for everything that is not explicitly bound/assigned to another zone, if FirewallD is being used, it is important for the default zone to set

Audit:

Run the following script to verify that **IF** FirewallD is being used, the default zone is set and follows local site policy:

```
#!/usr/bin/env bash

{
  l_output="" l_output2="" l_zone=""
  if systemctl is-enabled firewalld.service | grep -q 'enabled'; then
    l_zone="$(firewall-cmd --get-default-zone)"
    if [ -n "$l_zone" ]; then
      l_output=" - The default zone is set to: \"$l_zone\""
    else
      l_output2=" - The default zone is not set"
    fi
  else
    l_output=" - FirewallD is not in use on the system"
  fi
  if [ -z "$l_output2" ]; then
    echo -e "\n- Audit Results:\n ** Pass **\n$l_output\n"
  else
    echo -e "\n- Audit Results:\n ** Fail **\n$l_output2\n"
  fi
}
```


Remediation:

Run the following script to set the default zone:

```
#!/usr/bin/env bash

{
  l_zone="public" # <- Update to local site zone name if desired
  l_zone=""
  if systemctl is-enabled firewalld.service | grep -q 'enabled'; then
    l_zone="$(firewall-cmd --get-default-zone)"
    if [ "$l_zone" = "$l_zone" ]; then
      echo -e "\n - The default zone is set to: \"$l_zone\"\n - No remediation required"
    elif [ -n "$l_zone" ]; then
      echo -e "\n - The default zone is set to: \"$l_zone\"\n - Updating default zone to: \"$l_zone\""
      firewall-cmd --set-default-zone="$l_zone"
    else
      echo -e "\n - The default zone is set to: \"$l_zone\"\n - Updating default zone to: \"$l_zone\""
      firewall-cmd --set-default-zone="$l_zone"
    fi
  else
    echo -e "\n - FirewallD is not in use on the system\n - No remediation required"
  fi
}
```

Default Value:

public

References:

1. <https://firewalld.org/documentation>
2. <https://firewalld.org/documentation/man-pages/firewalld.zone>
3. NIST SP 800-53 Rev. 5: CA-9, SC-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.4 Implement and Manage a Firewall on Servers Implement and manage a firewall on servers, where supported. Example implementations include a virtual firewall, operating system firewall, or a third-party firewall agent.			
v7	9.4 Apply Host-based Firewalls or Port Filtering Apply host-based firewalls or port filtering tools on end systems, with a default-deny rule that drops all traffic except those services and ports that are explicitly allowed.			

3.4.2.2 Ensure at least one nftables table exists (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Tables hold chains. Each table only has one address family and only applies to packets of this family. Tables can have one of five families.

Rationale:

Without a table, nftables will not filter network traffic.

Impact:

Adding or modifying firewall rules can cause loss of connectivity to the system

Audit:

Run the following command to verify that a nftables table exists:

```
# nft list tables
```

Return should include a list of nftables. At least one table should be returned:

Example from a system using FirewallD:

```
table inet firewallld
```

Remediation:

Run the following command to create a table in nftables

```
# nft create table inet <table name>
```

Example if FirewallD is not in use on the system:

```
# nft create table inet filter
```

Note: FirewallD uses the `table inet firewallld` `NFTables` table that is created when FirewallD is installed.

References:

1. NIST SP 800-53 Rev. 5: CA-9

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.4 Implement and Manage a Firewall on Servers Implement and manage a firewall on servers, where supported. Example implementations include a virtual firewall, operating system firewall, or a third-party firewall agent.			
v7	9.4 Apply Host-based Firewalls or Port Filtering Apply host-based firewalls or port filtering tools on end systems, with a default-deny rule that drops all traffic except those services and ports that are explicitly allowed.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.004	TA0011	M1031, M1037

3.4.2.3 Ensure nftables base chains exist (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Chains are containers for rules. They exist in two kinds, base chains and regular chains. A base chain is an entry point for packets from the networking stack, a regular chain may be used as jump target and is used for better rule organization.

Rationale:

If a base chain doesn't exist with a hook for input, forward, and delete, packets that would flow through those chains will not be touched by nftables.

Impact:

If configuring over ssh, creating a base chain with a policy of drop will cause loss of connectivity.

Ensure that a rule allowing ssh has been added to the base chain prior to setting the base chain's policy to drop

Audit:

Run the following commands and verify that base chains exist for INPUT, FORWARD, and OUTPUT.

```
# nft list ruleset | grep 'hook input'
type filter hook input priority 0;

# nft list ruleset | grep 'hook forward'
type filter hook forward priority 0;

# nft list ruleset | grep 'hook output'
type filter hook output priority 0;
```

Note: When using FirewallD the base chains are installed by default

Remediation:

Run the following command to create the base chains:

```
# nft create chain inet <table name> <base chain name> { type filter hook  
<(input|forward|output)> priority 0 \; }
```

Example:

```
# nft create chain inet filter input { type filter hook input priority 0 \; }  
# nft create chain inet filter forward { type filter hook forward priority 0  
\; }  
# nft create chain inet filter output { type filter hook output priority 0 \;  
}
```

References:

1. NIST SP 800-53 Rev. 5: CA-9, SC-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.4 Implement and Manage a Firewall on Servers Implement and manage a firewall on servers, where supported. Example implementations include a virtual firewall, operating system firewall, or a third-party firewall agent.			
v7	9.4 Apply Host-based Firewalls or Port Filtering Apply host-based firewalls or port filtering tools on end systems, with a default-deny rule that drops all traffic except those services and ports that are explicitly allowed.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.004	TA0005	

3.4.2.4 Ensure host based firewall loopback traffic is configured (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Configure the loopback interface to accept traffic. Configure all other interfaces to deny traffic to the loopback network

Rationale:

Loopback traffic is generated between processes on machine and is typically critical to operation of the system. The loopback interface is the only place that loopback network traffic should be seen, all other interfaces should ignore traffic on this network as an anti-spoofing measure.

Audit:

Run the following script to verify that the loopback interface is configured:

```
#!/usr/bin/env bash

{
    l_output="" l_output2=""
    if nft list ruleset | awk '/hook\s+input\s+/,/\}\s*(#.*)?$/ ' | grep -Pq --
'\H+\h+"lo"\h+accept'; then
        l_output="$l_output\n - Network traffic to the loopback address is
correctly set to accept"
    else
        l_output2="$l_output2\n - Network traffic to the loopback address is
not set to accept"
    fi
    l_ipsaddr="$(nft list ruleset | awk
'/filter_IN_public_deny|hook\s+input\s+/,/\}\s*(#.*)?$/ ' | grep -P --
'ip\h+saddr')"
    if grep -Pq --
'ip\h+saddr\h+127\.0\.0\.0\/8\h+(counter\h+packets\h+\d+\h+bytes\h+\d+\h+)?dr
op' <<< "$l_ipsaddr" || grep -Pq --
'ip\h+daddr\h+!\=\h+127\.0\.0\.1\h+ip\h+saddr\h+127\.0\.0\.1\h+drop' <<<
"$l_ipsaddr"; then
        l_output="$l_output\n - IPv4 network traffic from loopback address
correctly set to drop"
    else
        l_output2="$l_output2\n - IPv4 network traffic from loopback address
not set to drop"
    fi
    if grep -Pq -- '^h*0h*$' /sys/module/ipv6/parameters/disable; then
        l_ip6saddr="$(nft list ruleset | awk '/filter_IN_public_deny|hook
input/,/}/ ' | grep 'ip6 saddr')"
        if grep -Pq
'ip6\h+saddr\h+::1\h+(counter\h+packets\h+\d+\h+bytes\h+\d+\h+)?drop' <<<
"$l_ip6saddr" || grep -Pq --
'ip6\h+daddr\h+!\=\h+::1\h+ip6\h+saddr\h+::1\h+drop' <<< "$l_ip6saddr"; then
            l_output="$l_output\n - IPv6 network traffic from loopback address
correctly set to drop"
        else
            l_output2="$l_output2\n - IPv6 network traffic from loopback address
not set to drop"
        fi
    fi
    if [ -z "$l_output2" ]; then
        echo -e "\n- Audit Result:\n    *** PASS ***\n$l_output"
    else
        echo -e "\n- Audit Result:\n    *** FAIL ***\n$l_output2\n\n - Correctly
set:\n$l_output"
    fi
}
```

Remediation:

Run the following script to implement the loopback rules:


```
#!/usr/bin/env bash

{
    l_hbfbw=""
    if systemctl is-enabled firewalld.service | grep -q 'enabled' && systemctl is-enabled nftables.service | grep -q 'enabled'; then
        echo -e "\n - Error - Both FirewallD and NFTables are enabled\n - Please follow recommendation: \"Ensure a single firewall configuration utility is in use\""
    elif ! systemctl is-enabled firewalld.service | grep -q 'enabled' && ! systemctl is-enabled nftables.service | grep -q 'enabled'; then
        echo -e "\n - Error - Neither FirewallD or NFTables is enabled\n - Please follow recommendation: \"Ensure a single firewall configuration utility is in use\""
    else
        if systemctl is-enabled firewalld.service | grep -q 'enabled' && ! systemctl is-enabled nftables.service | grep -q 'enabled'; then
            echo -e "\n - FirewallD is in use on the system" && l_hbfbw="fwd"
        elif ! systemctl is-enabled firewalld.service | grep -q 'enabled' && systemctl is-enabled nftables.service | grep -q 'enabled'; then
            echo -e "\n - NFTables is in use on the system" && l_hbfbw="nft"
        fi
        l_ipsaddr="$(nft list ruleset | awk '/filter_IN_public_deny|hook\s+input\s+/,/\)\s*(#.*)?$/ ' | grep -P -- 'ip\h+saddr')"
```

```
        if ! nft list ruleset | awk '/hook\s+input\s+/,/\)\s*(#.*)?$/ ' | grep -Pq -- '\H+\h+lo"\h+accept'; then
            echo -e "\n - Enabling input to accept for loopback address"
            if [ "$l_hbfbw" = "fwd" ]; then
                firewall-cmd --permanent --zone=trusted --add-interface=lo
                firewall-cmd --reload
            elif [ "$l_hbfbw" = "nft" ]; then
                nft add rule inet filter input iif lo accept
            fi
        fi
        if ! grep -Pq -- 'ip\h+saddr\h+127\.0\.0\.0\8\h+(counter\h+packets\h+d\h+bytes\h+d\h+)?drop' <<< "$l_ipsaddr" && ! grep -Pq -- 'ip\h+daddr\h+!\h+=\h+127\.0\.0\.1\h+ip\h+saddr\h+127\.0\.0\.1\h+drop' <<< "$l_ipsaddr"; then
            echo -e "\n - Setting IPv4 network traffic from loopback address to drop"
            if [ "$l_hbfbw" = "fwd" ]; then
                firewall-cmd --permanent --add-rich-rule='rule family=ipv4 source address="127.0.0.1" destination not address="127.0.0.1" drop'
                firewall-cmd --permanent --zone=trusted --add-rich-rule='rule family=ipv4 source address="127.0.0.1" destination not address="127.0.0.1" drop'
                firewall-cmd --reload
            elif [ "$l_hbfbw" = "nft" ]; then
                nft create rule inet filter input ip saddr 127.0.0.0/8 counter drop
            fi
        fi
        if grep -Pq -- '^h*0h*$' /sys/module/ipv6/parameters/disable; then
            l_ip6saddr="$(nft list ruleset | awk '/filter_IN_public_deny|hook input/,/)/ ' | grep 'ip6 saddr')"
```

```
            if ! grep -Pq 'ip6\h+saddr\h+::1\h+(counter\h+packets\h+d\h+bytes\h+d\h+)?drop' <<< "$l_ip6saddr" && ! grep -Pq -- 'ip6\h+daddr\h+!\h+=:1\h+ip6\h+saddr\h+::1\h+drop' <<< "$l_ip6saddr"; then
                echo -e "\n - Setting IPv6 network traffic from loopback address to drop"
                if [ "$l_hbfbw" = "fwd" ]; then
                    firewall-cmd --permanent --add-rich-rule='rule family=ipv6 source address="::1" destination not address="::1" drop'
                    firewall-cmd --permanent --zone=trusted --add-rich-rule='rule family=ipv6 source address="::1" destination not address="::1" drop'
                    firewall-cmd --reload
                elif [ "$l_hbfbw" = "nft" ]; then
                    nft add rule inet filter input ip6 saddr ::1 counter drop
                fi
            fi
        fi
    fi
}
```

References:

1. NIST SP 800-53 Rev. 5: CA-9

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.4 Implement and Manage a Firewall on Servers Implement and manage a firewall on servers, where supported. Example implementations include a virtual firewall, operating system firewall, or a third-party firewall agent.			
v7	9.4 Apply Host-based Firewalls or Port Filtering Apply host-based firewalls or port filtering tools on end systems, with a default-deny rule that drops all traffic except those services and ports that are explicitly allowed.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.004	TA0005	

3.4.2.5 Ensure firewalld drops unnecessary services and ports (Manual)

Profile Applicability:

- Level 2 - Server

Description:

Services and ports can be accepted or explicitly rejected or dropped by a zone.

For every zone, you can set a default behavior that handles incoming traffic that is not further specified. Such behavior is defined by setting the target of the zone. There are three options - default, ACCEPT, REJECT, and DROP.

- ACCEPT - you accept all incoming packets except those disabled by a specific rule.
- REJECT - you disable all incoming packets except those that you have allowed in specific rules and the source machine is informed about the rejection.
- DROP - you disable all incoming packets except those that you have allowed in specific rules and no information sent to the source machine.

Rationale:

To reduce the attack surface of a system, all services and ports should be blocked unless required

Audit:

Run the following command and review output to ensure that listed services and ports follow site policy.

```
# systemctl is-enabled firewalld.service | grep -q 'enabled' && firewall-cmd --list-all --zone="$(firewall-cmd --list-all | awk '/\ (active\)/ { print $1 }')" | grep -P -- '^h*(services:|ports:)'
```

Remediation:

If Firewalld is in use on the system:

Run the following command to remove an unnecessary service:

```
# firewall-cmd --remove-service=<service>
```

Example:

```
# firewall-cmd --remove-service=cockpit
```

Run the following command to remove an unnecessary port:

```
# firewall-cmd --remove-port=<port-number>/<port-type>
```

Example:

```
# firewall-cmd --remove-port=25/tcp
```

Run the following command to make new settings persistent:

```
# firewall-cmd --runtime-to-permanent
```

References:

1. firewalld.service(5)
2. https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html/securing_networks/using-and-configuring-firewalls_securing-networks
3. NIST SP 800-53 Rev. 5: CA-9

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.4 Implement and Manage a Firewall on Servers Implement and manage a firewall on servers, where supported. Example implementations include a virtual firewall, operating system firewall, or a third-party firewall agent.			
v7	9.4 Apply Host-based Firewalls or Port Filtering Apply host-based firewalls or port filtering tools on end systems, with a default-deny rule that drops all traffic except those services and ports that are explicitly allowed.			

3.4.2.6 Ensure nftables established connections are configured (Manual)

Profile Applicability:

- Level 2 - Server

Description:

Configure the firewall rules for new outbound and established connections

Rationale:

If rules are not in place for established connections, all packets will be dropped by the default policy preventing network usage.

Audit:

If NFTables utility is in use on your system:

Run the following commands and verify all rules for established incoming connections match site policy:

```
# systemctl is-enabled nftables.service | grep -q 'enabled' && nft list ruleset | awk '/hook input/,/}/' | grep 'ct state'
```

Output should be similar to:

```
ip protocol tcp ct state established accept
ip protocol udp ct state established accept
ip protocol icmp ct state established accept
```

Remediation:

If NFTables utility is in use on your system:

Configure nftables in accordance with site policy. The following commands will implement a policy to allow all established connections:

```
# systemctl is-enabled nftables.service | grep -q 'enabled' && nft add rule inet filter input ip protocol tcp ct state established accept
# systemctl is-enabled nftables.service | grep -q 'enabled' && nft add rule inet filter input ip protocol udp ct state established accept
# systemctl is-enabled nftables.service | grep -q 'enabled' && nft add rule inet filter input ip protocol icmp ct state established accept
```

References:

1. NIST SP 800-53 Rev. 5: CA-9

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.4 <u>Implement and Manage a Firewall on Servers</u> Implement and manage a firewall on servers, where supported. Example implementations include a virtual firewall, operating system firewall, or a third-party firewall agent.			
v7	9.4 <u>Apply Host-based Firewalls or Port Filtering</u> Apply host-based firewalls or port filtering tools on end systems, with a default-deny rule that drops all traffic except those services and ports that are explicitly allowed.			

3.4.2.7 Ensure nftables default deny firewall policy (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Base chain policy is the default verdict that will be applied to packets reaching the end of the chain.

Rationale:

There are two policies: accept (Default) and drop. If the policy is set to `accept`, the firewall will accept any packet that is not configured to be denied and the packet will continue traversing the network stack.

It is easier to explicitly permit acceptable usage than to deny unacceptable usage.

Note: Changing firewall settings while connected over the network can result in being locked out of the system.

Impact:

If configuring nftables over ssh, creating a base chain with a policy of drop will cause loss of connectivity.

Ensure that a rule allowing `ssh` has been added to the base chain prior to setting the base chain's policy to drop

Audit:

If Nftables utility is in use on your system:

Run the following commands and verify that base chains contain a policy of `DROP`.

```
# systemctl --quiet is-enabled nftables.service && nft list ruleset | grep 'hook input' | grep -v 'policy drop'
```

Nothing should be returned

```
# systemctl --quiet is-enabled nftables.service && nft list ruleset | grep 'hook forward' | grep -v 'policy drop'
```

Nothing should be returned

Remediation:

If NfTables utility is in use on your system:

Run the following command for the base chains with the input, forward, and output hooks to implement a default DROP policy:

```
# nft chain <table family> <table name> <chain name> { policy drop \; }
```

Example:

```
# nft chain inet filter input { policy drop \; }  
# nft chain inet filter forward { policy drop \; }
```

Default Value:

accept

References:

1. Manual Page nft
2. NIST SP 800-53 Rev. 5: CA-9

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.4 <u>Implement and Manage a Firewall on Servers</u> Implement and manage a firewall on servers, where supported. Example implementations include a virtual firewall, operating system firewall, or a third-party firewall agent.			
v7	9.4 <u>Apply Host-based Firewalls or Port Filtering</u> Apply host-based firewalls or port filtering tools on end systems, with a default-deny rule that drops all traffic except those services and ports that are explicitly allowed.			

4 Access, Authentication and Authorization

4.1 Configure time-based job schedulers

`cron` is a time-based job scheduler used to schedule jobs, commands or shell scripts, to run periodically at fixed times, dates, or intervals.

`at` provides the ability to execute a command or shell script at a specified date and hour, or after a given interval of time.

Other methods exist for scheduling jobs, such as `systemd timers`. If another method is used, it should be secured in accordance with local site policy

Note: `systemd timers` are `systemd` unit files whose name ends in `.timer` that control `.service` files or events. Timers can be used as an alternative to `cron` and `at`. Timers have built-in support for calendar time events, monotonic time events, and can be run asynchronously

If `cron` and `at` are not installed, this section can be skipped.

4.1.1 Ensure cron daemon is installed and enabled (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `cron` daemon is used to execute batch jobs on the system.

Rationale:

While there may not be user jobs that need to be run on the system, the system does have maintenance jobs that may include security monitoring that have to run, and `cron` is used to execute them.

Audit:

-- IF -- Cron is being used to execute batch jobs on the system:

Run the following commands to verify `cron` is installed and enabled:

```
# rpm -q cronie
cronie-1.5.7-1.amzn2023.0.2.x86_64
# systemctl is-enabled crond
enabled
```

Verify result is "enabled".

Remediation:

-- IF -- Cron is being used to execute batch jobs on the system:

Run the following command to install and enable `cron`:

```
# dnf install cronie
# systemctl --now enable crond
```

References:

1. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

Additional Information:

Additional methods of enabling a service exist. Consult your distribution documentation for appropriate methods.

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.001	TA0005	M1018

4.1.2 Ensure permissions on /etc/crontab are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `/etc/crontab` file is used by `cron` to control its own jobs. The commands in this item make sure that `root` is the user and group owner of the file and that only the owner can access the file.

Rationale:

This file contains information on what system jobs are run by `cron`. Write access to this file could provide unprivileged users with the ability to elevate their privileges. Read access to this file could provide users with the ability to gain insight on system jobs that run on the system and could provide them a way to gain unauthorized privileged access.

Audit:

-- IF -- Cron is being used to execute batch jobs on the system

Run the following command and verify `Uid` and `Gid` are both `0/root` and `Access` does not grant permissions to `group` or `other` :

```
# stat /etc/crontab
Access: (0600/-rw-----)  Uid: (    0/   root)   Gid: (    0/   root)
```

Remediation:

-- IF -- Cron is being used to execute batch jobs on the system

Run the following commands to set ownership and permissions on `/etc/crontab`:

```
# chown root:root /etc/crontab
# chmod og-rwx /etc/crontab
```

References:

1. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1053, T1053.003	TA0002, TA0007	M1018

4.1.3 Ensure permissions on /etc/cron.hourly are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

This directory contains system `cron` jobs that need to run on an hourly basis. The files in this directory cannot be manipulated by the `crontab` command, but are instead edited by system administrators using a text editor. The commands below restrict read/write and search access to user and group root, preventing regular users from accessing this directory.

Rationale:

Granting write access to this directory for non-privileged users could provide them the means for gaining unauthorized elevated privileges. Granting read access to this directory could give an unprivileged user insight in how to gain elevated privileges or circumvent auditing controls.

Audit:

-- IF -- Cron is being used to execute batch jobs on the system:

Run the following command and verify `Uid` and `Gid` are both `0/root` and `Access` does not grant permissions to `group` or `other` :

```
# stat /etc/cron.hourly
Access: (0700/drwx-----)  Uid: (    0/    root)  Gid: (    0/    root)
```

Remediation:

-- IF -- Cron is being used to execute batch jobs on the system:

Run the following commands to set ownership and permissions on `/etc/cron.hourly` :

```
# chown root:root /etc/cron.hourly
# chmod og-rwx /etc/cron.hourly
```

References:

1. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1053, T1053.003	TA0002, TA0007	M1018

4.1.4 Ensure permissions on /etc/cron.daily are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `/etc/cron.daily` directory contains system cron jobs that need to run on a daily basis. The files in this directory cannot be manipulated by the `crontab` command, but are instead edited by system administrators using a text editor. The commands below restrict read/write and search access to user and group root, preventing regular users from accessing this directory.

Rationale:

Granting write access to this directory for non-privileged users could provide them the means for gaining unauthorized elevated privileges. Granting read access to this directory could give an unprivileged user insight in how to gain elevated privileges or circumvent auditing controls.

Audit:

-- IF -- Cron is being used to execute batch jobs on the system:

Run the following command and verify `Uid` and `Gid` are both `0/root` and `Access` does not grant permissions to `group` or `other` :

```
# stat /etc/cron.daily
Access: (0700/drwx-----)  Uid: (    0/    root)  Gid: (    0/    root)
```

Remediation:

-- IF -- Cron is being used to execute batch jobs on the system:

Run the following commands to set ownership and permissions on `/etc/cron.daily` :

```
# chown root:root /etc/cron.daily
# chmod og-rwx /etc/cron.daily
```

References:

1. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1053, T1053.003	TA0002, TA0007	M1018

4.1.5 Ensure permissions on /etc/cron.weekly are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `/etc/cron.weekly` directory contains system cron jobs that need to run on a weekly basis. The files in this directory cannot be manipulated by the `crontab` command, but are instead edited by system administrators using a text editor. The commands below restrict read/write and search access to user and group root, preventing regular users from accessing this directory.

Rationale:

Granting write access to this directory for non-privileged users could provide them the means for gaining unauthorized elevated privileges. Granting read access to this directory could give an unprivileged user insight in how to gain elevated privileges or circumvent auditing controls.

Audit:

-- IF -- Cron is being used to execute batch jobs on the system:

Run the following command and verify `Uid` and `Gid` are both `0/root` and `Access` does not grant permissions to `group` or `other` :

```
# stat /etc/cron.weekly
Access: (0700/drwx-----)  Uid: (    0/    root)  Gid: (    0/    root)
```

Remediation:

-- IF -- Cron is being used to execute batch jobs on the system:

Run the following commands to set ownership and permissions on `/etc/cron.weekly` :

```
# chown root:root /etc/cron.weekly
# chmod og-rwx /etc/cron.weekly
```

References:

1. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1053, T1053.003	TA0002, TA0007	M1018

4.1.6 Ensure permissions on /etc/cron.monthly are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `/etc/cron.monthly` directory contains system cron jobs that need to run on a monthly basis. The files in this directory cannot be manipulated by the `crontab` command, but are instead edited by system administrators using a text editor. The commands below restrict read/write and search access to user and group root, preventing regular users from accessing this directory.

Rationale:

Granting write access to this directory for non-privileged users could provide them the means for gaining unauthorized elevated privileges. Granting read access to this directory could give an unprivileged user insight in how to gain elevated privileges or circumvent auditing controls.

Audit:

-- IF -- Cron is being used to execute batch jobs on the system:

Run the following command and verify `Uid` and `Gid` are both `0/root` and `Access` does not grant permissions to `group` or `other`:

```
# stat /etc/cron.monthly
Access: (0700/drwx-----)  Uid: (    0/    root)  Gid: (    0/    root)
```

Remediation:

-- IF -- Cron is being used to execute batch jobs on the system:

Run the following commands to set ownership and permissions on `/etc/cron.monthly`:

```
# chown root:root /etc/cron.monthly
# chmod og-rwx /etc/cron.monthly
```

References:

1. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1053, T1053.003	TA0002, TA0007	M1018

4.1.7 Ensure permissions on /etc/cron.d are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `/etc/cron.d` directory contains system `cron` jobs that need to run in a similar manner to the hourly, daily, weekly and monthly jobs from `/etc/crontab`, but require more granular control as to when they run. The files in this directory cannot be manipulated by the `crontab` command, but are instead edited by system administrators using a text editor. The commands below restrict read/write and search access to user and group root, preventing regular users from accessing this directory.

Rationale:

Granting write access to this directory for non-privileged users could provide them the means for gaining unauthorized elevated privileges. Granting read access to this directory could give an unprivileged user insight in how to gain elevated privileges or circumvent auditing controls.

Audit:

-- IF -- Cron is being used to execute batch jobs on the system:

Run the following command and verify `Uid` and `Gid` are both `0/root` and `Access` does not grant permissions to `group` or `other`:

```
# stat /etc/cron.d
Access: (0700/drwx-----)  Uid: (    0/    root)  Gid: (    0/    root)
```

Remediation:

-- IF -- Cron is being used to execute batch jobs on the system:

Run the following commands to set ownership and permissions on `/etc/cron.d`:

```
# chown root:root /etc/cron.d
# chmod og-rwx /etc/cron.d
```

References:

1. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1053, T1053.003	TA0002, TA0007	M1022

4.1.8 Ensure cron is restricted to authorized users (Automated)

Profile Applicability:

- Level 1 - Server

Description:

If `cron` is installed in the system, configure `/etc/cron.allow` to allow specific users to use these services. If `/etc/cron.allow` does not exist, then `/etc/cron.deny` is checked. Any user not specifically defined in those files is allowed to use cron. By removing the file, only users in `/etc/cron.allow` are allowed to use cron.

Note: Even though a given user is not listed in `cron.allow`, cron jobs can still be run as that user. The `cron.allow` file only controls administrative access to the `crontab` command for scheduling and modifying cron jobs.

Rationale:

On many systems, only the system administrator is authorized to schedule `cron` jobs. Using the `cron.allow` file to control who can run `cron` jobs enforces this policy. It is easier to manage an allow list than a deny list. In a deny list, you could potentially add a user ID to the system and forget to add it to the deny files.

Audit:

-- IF -- Cron is being used to execute batch jobs on the system:
Run the following script:

```
#!/usr/bin/env bash

{
    if rpm -q crone >/dev/null; then
        [ -e /etc/cron.deny ] && echo "Fail: cron.deny exists"
        if [ ! -e /etc/cron.allow ]; then
            echo "Fail: cron.allow doesn't exist"
        else
            ! stat -Lc "%a" /etc/cron.allow | grep -Eq "[0,2,4,6]00" && echo
            "Fail: cron.allow mode too permissive"
            ! stat -Lc "%u:%g" /etc/cron.allow | grep -Eq "^0:0$" && echo "Fail:
cron.allow owner and/or group not root"
        fi
        if [ ! -e /etc/cron.deny ] && [ -e /etc/cron.allow ] && stat -Lc "%a"
/etc/cron.allow | grep -Eq "[0,2,4,6]00" \
            && stat -Lc "%u:%g" /etc/cron.allow | grep -Eq "^0:0$"; then
            echo "Pass"
        fi
    else
        echo "Pass: cron is not installed on the system"
    fi
}
```

Verify the output of the script includes `Pass`

Remediation:

-- IF -- Cron is being used to execute batch jobs on the system:

Run the following script to remove `/etc/cron.deny`, create `/etc/cron.allow`, and set the file mode on `/etc/cron.allow`:

```
#!/usr/bin/env bash

{
  if rpm -q crontab >/dev/null; then
    [ -e /etc/cron.deny ] && rm -f /etc/cron.deny
    [ ! -e /etc/cron.allow ] && touch /etc/cron.allow
    chown root:root /etc/cron.allow
    chmod u-x,go-rwx /etc/cron.allow
  else
    echo "cron is not installed on the system"
  fi
}
```

OR Run the following command to remove `crontab`:

```
# dnf remove crontab
```

References:

1. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 Configure Data Access Control Lists Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 Protect Information through Access Control Lists Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1053, T1053.003	TA0002, TA0007	M1022

4.1.9 Ensure at is restricted to authorized users (Automated)

Profile Applicability:

- Level 1 - Server

Description:

-- IF -- `at` is installed in the system:

Configure `/etc/at.allow` to allow specific users to use these services. If `/etc/at.allow` does not exist, then `/etc/at.deny` is checked. Any user not specifically defined in those files is allowed to use `at`. By removing the file, only users in `/etc/at.allow` are allowed to use `at`.

Note: Even though a given user is not listed in `at.allow`, `at` jobs can still be run as that user. The `at.allow` file only controls administrative access to the `at` command for scheduling and modifying `at` jobs.

Rationale:

On many systems, only the system administrator is authorized to schedule `at` jobs. Using the `at.allow` file to control who can run `at` jobs enforces this policy. It is easier to manage an allow list than a deny list. In a deny list, you could potentially add a user ID to the system and forget to add it to the deny files.

Audit:

Run the following script:

```
#!/usr/bin/env bash

{
    if rpm -q at >/dev/null; then
        [ -e /etc/at.deny ] && echo "Fail: at.deny exists"
        if [ ! -e /etc/at.allow ]; then
            echo "Fail: at.allow doesn't exist"
        else
            ! stat -Lc "%a" /etc/at.allow | grep -Eq "[0,2,4,6]00" && echo
            "Fail: at.allow mode too permissive"
            ! stat -Lc "%u:%g" /etc/at.allow | grep -Eq "^0:0$" && echo "Fail:
at.allow owner and/or group not root"
        fi
        if [ ! -e /etc/at.deny ] && [ -e /etc/at.allow ] && stat -Lc "%a"
/etc/at.allow | grep -Eq "[0,2,4,6]00" \
            && stat -Lc "%u:%g" /etc/at.allow | grep -Eq "^0:0$"; then
            echo "Pass"
        fi
    else
        echo "Pass: at is not installed on the system"
    fi
}
```

Verify the output of the script includes `Pass`

Remediation:

Run the following script to remove `/etc/at.deny`, create `/etc/at.allow`, and set the file mode for `/etc/at.allow`:

```
#!/usr/bin/env bash

{
    if rpm -q at >/dev/null; then
        [ -e /etc/at.deny ] && rm -f /etc/at.deny
        [ ! -e /etc/at.allow ] && touch /etc/at.allow
        chown root:root /etc/at.allow
        chmod u-x,go-rwx /etc/at.allow
    else
        echo "at is not installed on the system"
    fi
}
```

OR Run the following command to remove `at`:

```
# dnf remove at
```

References:

1. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1053, T1053.003	TA0002, TA0007	M1026

4.2 Configure SSH Server

SSH is a secure, encrypted replacement for common login services such as `telnet`, `ftp`, `rlogin`, `rsh`, and `rcp`. It is strongly recommended that sites abandon older clear-text login protocols and use SSH to prevent session hijacking and sniffing of sensitive data off the network.

Note:

- The recommendations in this section only apply if the SSH daemon is installed on the system, if remote access is **not** required the SSH daemon can be removed and this section skipped.
- By default the `/etc/ssh/sshd_config` file includes a `Include /etc/ssh/sshd_config.d/*.conf` entry.
 - This entry will automatically include settings in `*.conf` files in the `/etc/ssh/sshd_config.d/` directory to be included in the `sshd` configuration.
 - `Include:`
 - Include the specified configuration file(s).
 - Multiple pathnames may be specified and each pathname may contain `glob(7)` wildcards that will be expanded and processed in lexical order.
 - Files without absolute paths are assumed to be in `/etc/ssh`.
 - An `Include` directive may appear inside a `Match` block to perform conditional inclusion.
- If additional files are added to the `Include` parameter, these files must also be accounted for in the audit and remediation procedures in this section.
- Once all configuration changes have been made, the `sshd` configuration must be reloaded:

Command to re-load the SSH daemon configuration:

```
# systemctl reload sshd
```

Command to remove the SSH daemon:

```
# dnf remove openssh-server
```

4.2.1 Ensure permissions on /etc/ssh/sshd_config are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `/etc/ssh/sshd_config` file contains configuration specifications for `sshd`. The command below sets the owner and group of the file to root.

Rationale:

The `/etc/ssh/sshd_config` file needs to be protected from unauthorized changes by non-privileged users.

Audit:

Run the following command and verify `Uid` and `Gid` are both `0/root` and `Access` does not grant permissions to `group` or `other`:

```
# stat -Lc "%n %a %u/%U %g/%G" /etc/ssh/sshd_config
/etc/ssh/sshd_config 600 0/root 0/root
```

Remediation:

Run the following commands to set ownership and permissions on `/etc/ssh/sshd_config`:

```
# chown root:root /etc/ssh/sshd_config
# chmod u-x,go-rwx /etc/ssh/sshd_config
```

Default Value:

`/etc/ssh/sshd_config 600 0/root 0/root`

References:

1. NIST SP 800-53 Rev. 5: AC-3. MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1098, T1098.004, T1543, T1543.002	TA0005	M1022

4.2.2 Ensure permissions on SSH private host key files are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

An SSH private key is one of two files used in SSH public key authentication. In this authentication method, the possession of the private key is proof of identity. Only a private key that corresponds to a public key will be able to authenticate successfully. The private keys need to be stored and handled carefully, and no copies of the private key should be distributed.

Rationale:

If an unauthorized user obtains the private SSH host key file, the host could be impersonated

Audit:

Run the following script to verify SSH private host key files are mode 0600 or more restrictive, owned by the root user, and owned by the group root or group designated to own openSSH private keys:

```

#!/usr/bin/env bash

{
    l_output="" l_output2=""
    l_skgn="ssh_keys" # Group designated to own openSSH keys
    l_skgid="$(awk -F: '($1 == "'$l_skgn"') {print $3}' /etc/group)" # Get
gid of group
    [ -n "$l_skgid" ] && l_agroup="(root|$l_skgn)" || l_agroup="root"
    unset a_skarr && a_skarr=() # Clear and initialize array
    while IFS= read -r -d $'\0' l_file; do # Loop to populate array
        if grep -Pq ':\h+OpenSSH\h+private\h+key\b' <<< "$(file "$l_file")";
then
            a_skarr+=("$(stat -Lc '%n^%#a^%U^%G^%g' "$l_file")")
        fi
        done <<(find -L /etc/ssh -xdev -type f -print0)
        while IFS="^" read -r l_file l_mode l_owner l_group l_gid; do
            echo "File: \"$l_file\" Mode: \"$l_mode\" Owner: \"$l_owner\" Group:
\"$l_group\" GID: \"$l_gid\""
            l_out2=""
            [ "$l_gid" = "$l_skgid" ] && l_pmask="0137" || l_pmask="0177"
            l_maxperm="$( printf '%o' $(( 0777 & ~$l_pmask )) )"
            if [ $(( $l_mode & $l_pmask )) -gt 0 ]; then
                l_out2="$l_out2\n - Mode: \"$l_mode\" should be mode:
\"$l_maxperm\" or more restrictive"
            fi
            if [ "$l_owner" != "root" ]; then
                l_out2="$l_out2\n - Owned by: \"$l_owner\" should be owned by
\"root\""
            fi
            if [[ ! "$l_group" =~ $l_agroup ]]; then
                l_out2="$l_out2\n - Owned by group \"$l_group\" should be group
owned by: \"${l_agroup//|/ or }\""
            fi
            if [ -n "$l_out2" ]; then
                l_output2="$l_output2\n - File: \"$l_file\"$l_out2"
            else
                l_output="$l_output\n - File: \"$l_file\"\n - Correct: mode
($l_mode), owner ($l_owner), and group owner ($l_group) configured"
            fi
            done <<< "$(printf '%s\n' "${a_skarr[@]}")"
            unset a_skarr
            if [ -z "$l_output2" ]; then
                echo -e "\n- Audit Result:\n *** PASS ***\n- * Correctly set *
:\n$l_output\n"
            else
                echo -e "\n- Audit Result:\n ** FAIL **\n- * Reasons for audit
failure * :\n$l_output2\n"
                [ -n "$l_output" ] && echo -e " - * Correctly set * :\n$l_output\n"
            fi
        }
}

```

Remediation:

Run the following script to set mode, ownership, and group on the private SSH host key files:

```
#!/usr/bin/env bash

{
    l_output="" l_output2=""
    l_skgn="ssh_keys" # Group designated to own openSSH keys
    l_skgid="$(awk -F: '($1 == "'$l_skgn'"){print $3}' /etc/group)" # Get
gid of group
    if [ -n "$l_skgid" ]; then
        l_agroup="(root|$l_skgn)" && l_sgroup="$l_skgn" && l_mfix="u-x,g-wx,o-
rwx"
    else
        l_agroup="root" && l_sgroup="root" && l_mfix="u-x,go-rwx"
    fi
    unset a_skarr && a_skarr=() # Clear and initialize array
    while IFS= read -r -d $'\0' l_file; do # Loop to populate array
        if grep -Pq ':\h+OpenSSH\h+private\h+key\b' <<< "$(file "$l_file")";
then
            a_skarr+=("$((stat -Lc '%n^%a^%U^%G^%g' "$l_file")"))
        fi
    done << (find -L /etc/ssh -xdev -type f -print0)
    while IFS="^" read -r l_file l_mode l_owner l_group l_gid; do
        l_out2=""
        [ "$l_gid" = "$l_skgid" ] && l_pmask="0137" || l_pmask="0177"
        l_maxperm="$( printf '%o' $(( 0777 & ~$l_pmask )) )"
        if [ $(( $l_mode & $l_pmask )) -gt 0 ]; then
            l_out2="$l_out2\n  - Mode: \"$l_mode\" should be mode:
\"$l_maxperm\" or more restrictive\n  - Revoking excess permissions"
            chmod "$l_mfix" "$l_file"
        fi
        if [ "$l_owner" != "root" ]; then
            l_out2="$l_out2\n  - Owned by: \"$l_owner\" should be owned by
\"root\"\n  - Changing ownership to \"root\""
            chown root "$l_file"
        fi
        if [[ ! "$l_group" =~ $l_agroup ]]; then
            l_out2="$l_out2\n  - Owned by group \"$l_group\" should be group
owned by: \"${l_agroup//|/ or }\"\n  - Changing group ownership to
\"$l_sgroup\""
            chgrp "$l_sgroup" "$l_file"
        fi
        [ -n "$l_out2" ] && l_output2="$l_output2\n - File: \"$l_file\"$l_out2"
    done <<< "$(printf '%s\n' "${a_skarr[@]}")"
    unset a_skarr
    if [ -z "$l_output2" ]; then
        echo -e "\n- No access changes required\n"
    else
        echo -e "\n- Remediation results:\n$l_output2\n"
    fi
}
```

References:

1. NIST SP 800-53 Rev. 5: AC-3. MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1552, T1552.004	TA0003, TA0006	M1022

4.2.3 Ensure permissions on SSH public host key files are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

An SSH public key is one of two files used in SSH public key authentication. In this authentication method, a public key is a key that can be used for verifying digital signatures generated using a corresponding private key. Only a public key that corresponds to a private key will be able to authenticate successfully.

Rationale:

If a public host key file is modified by an unauthorized user, the SSH service may be compromised.

Audit:

Run the following command and verify Access does not grant write or execute permissions to group or other for all returned files:

Run the following script to verify SSH public host key files are mode 0644 or more restrictive, owned by the root user, and owned by the root group:

```
#!/usr/bin/env bash

{
    l_output="" l_output2=""
    l_pmask="0133"
    awk '{print}' <<< "$(find -L /etc/ssh -xdev -type f -exec stat -Lc "%n %#a %U %G" {} +)" | (while read -r l_file l_mode l_owner l_group; do
        if file "$l_file" | grep -Pq ':\h+OpenSSH\h+(\h+)\h+key\b';
    then
        l_maxperm="$( printf '%o' $(( 0777 & ~$l_pmask )) )"
        if [ $(( $l_mode & $l_pmask )) -gt 0 ]; then
            l_output2="$l_output2\n - Public key file: \"$l_file\" is mode \"$l_mode\" should be mode: \"$l_maxperm\" or more restrictive"
        else
            l_output="$l_output\n - Public key file: \"$l_file\" is mode \"$l_mode\" should be mode: \"$l_maxperm\" or more restrictive"
        fi
        if [ "$l_owner" != "root" ]; then
            l_output2="$l_output2\n - Public key file: \"$l_file\" is owned by: \"$l_owner\" should be owned by \"root\""
        else
            l_output="$l_output\n - Public key file: \"$l_file\" is owned by: \"$l_owner\" should be owned by \"root\""
        fi
        if [ "$l_group" != "root" ]; then
            l_output2="$l_output2\n - Public key file: \"$l_file\" is owned by group \"$l_group\" should belong to group \"root\""
        else
            l_output="$l_output\n - Public key file: \"$l_file\" is owned by group \"$l_group\" should belong to group \"root\""
        fi
    fi
done
if [ -z "$l_output2" ]; then
    echo -e "\n- Audit Result:\n *** PASS ***\n$l_output"
else
    echo -e "\n- Audit Result:\n *** FAIL ***\n$l_output2\n\n - Correctly set:\n$l_output"
fi
}
```

Remediation:

Run the following script to set mode, ownership, and group on the public SSH host key files:

```
#!/usr/bin/env bash

{
    l_pmask="0133"
    l_maxperm="$( printf '%o' $(( 0777 & ~$l_pmask )) )"
    awk '{print}' <<< "$(find -L /etc/ssh -xdev -type f -exec stat -Lc "%n %#a
%U %G" {} +)" | (while read -r l_file l_mode l_owner l_group; do
        if file "$l_file" | grep -Pq ':\h+OpenSSH\h+(\h+\h+)?public\h+key\b';
    then
        echo -e " - Checking private key file: \"$l_file\""
        if [ $(( $l_mode & $l_pmask )) -gt 0 ]; then
            echo -e " - File: \"$l_file\" is mode \"$l_mode\" changing to
mode: \"$l_maxperm\""
            chmod u-x,go-wx "$l_file"
        fi
        if [ "$l_owner" != "root" ]; then
            echo -e " - File: \"$l_file\" is owned by: \"$l_owner\" changing
owner to \"root\""
            chown root "$l_file"
        fi
        if [ "$l_group" != "root" ]; then
            echo -e " - File: \"$l_file\" is owned by group \"$l_group\"
changing to group \"root\""
            chgrp "root" "$l_file"
        fi
    fi
done
)
```

Default Value:

644 0/root 0/root

References:

1. NIST SP 800-53 Rev. 5: AC-3. MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	5.1 <u>Establish Secure Configurations</u> Maintain documented, standard security configuration standards for all authorized operating systems and software.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1557, T1557.000	TA0003, TA0006	M1022

4.2.4 Ensure SSH access is limited (Automated)

Profile Applicability:

- Level 1 - Server

Description:

There are several options available to limit which users and group can access the system via SSH. It is recommended that at least one of the following options be leveraged:

- `AllowUsers:`
 - The `AllowUsers` variable gives the system administrator the option of allowing specific users to `ssh` into the system. The list consists of space separated user names. Numeric user IDs are not recognized with this variable. If a system administrator wants to restrict user access further by only allowing the allowed users to log in from a particular host, the entry can be specified in the form of `user@host`.
- `AllowGroups:`
 - The `AllowGroups` variable gives the system administrator the option of allowing specific groups of users to `ssh` into the system. The list consists of space separated group names. Numeric group IDs are not recognized with this variable.
- `DenyUsers:`
 - The `DenyUsers` variable gives the system administrator the option of denying specific users to `ssh` into the system. The list consists of space separated user names. Numeric user IDs are not recognized with this variable. If a system administrator wants to restrict user access further by specifically denying a user's access from a particular host, the entry can be specified in the form of `user@host`.
- `DenyGroups:`
 - The `DenyGroups` variable gives the system administrator the option of denying specific groups of users to `ssh` into the system. The list consists of space separated group names. Numeric group IDs are not recognized with this variable.

Rationale:

Restricting which users can remotely access the system via SSH will help ensure that only authorized users access the system.

Audit:

Run the following commands and verify the output:

```
# sshd -T -C user=root -C host="$(hostname)" -C addr="$(grep $(hostname)
/etc/hosts | awk '{print $1}')" | grep -Pi
'^\h*(allow|deny) (users|groups) \h+\H+(\h+.)? $'

# grep -Pi '^ \h*(allow|deny) (users|groups) \h+\H+(\h+.)? $'
/etc/ssh/sshd_config /etc/ssh/sshd_config.d/*.conf
```

Verify that the output of both commands matches at least one of the following lines:

```
allowusers <userlist>
allowgroups <grouplist>
denyusers <userlist>
denygroups <grouplist>
```

Remediation:

Edit or create a file ending in *.conf in the /etc/ssh/sshd_config.d/ directory or the /etc/ssh/sshd_config file and set one or more of the parameters as follows:

```
AllowUsers <userlist>
```

-OR-

```
AllowGroups <grouplist>
```

-OR-

```
DenyUsers <userlist>
```

-OR-

```
DenyGroups <grouplist>
```

Default Value:

None

References:

1. SSHD_CONFIG(5)
2. NIST SP 800-53 Rev. 5: AC-3. MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	4.3 <u>Ensure the Use of Dedicated Administrative Accounts</u> Ensure that all users with administrative account access use a dedicated or secondary account for elevated activities. This account should only be used for administrative activities and not internet browsing, email, or similar activities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1021, T1021.004	TA0008	M1018

4.2.5 Ensure SSH LogLevel is appropriate (Automated)

Profile Applicability:

- Level 1 - Server

Description:

`INFO` level is the basic level that only records login activity of SSH users. In many situations, such as Incident Response, it is important to determine when a particular user was active on a system. The logout record can eliminate those users who disconnected, which helps narrow the field.

`VERBOSE` level specifies that login and logout activity as well as the key fingerprint for any SSH key used for login will be logged. This information is important for SSH key management, especially in legacy environments.

Rationale:

SSH provides several logging levels with varying amounts of verbosity. `DEBUG` is specifically **not** recommended other than strictly for debugging SSH communications since it provides so much data that it is difficult to identify important security information.

Audit:

Run the following command and verify that output matches `loglevel VERBOSE` or `loglevel INFO`:

```
# sshd -T -C user=root -C host="$(hostname)" -C addr="$(grep $(hostname) /etc/hosts | awk '{print $1}')" | grep loglevel

loglevel VERBOSE
-or-
loglevel INFO
```

Run the following command and verify the output matches:

```
# grep -Pi -- '^h*loglevel' /etc/ssh/sshd_config
/etc/ssh/sshd_config.d/*.conf | grep -Evi '(VERBOSE|INFO) '

Nothing should be returned
```

Remediation:

Edit or create a file ending in *.conf in the /etc/ssh/sshd_config.d/ directory or the /etc/ssh/sshd_config file and set the LogLevel parameter as follows:

```
LogLevel VERBOSE
```

OR

```
LogLevel INFO
```

Run the following command to comment out any LogLevel parameter entries in files ending in *.conf in the /etc/ssh/sshd_config.d/ directory or the /etc/ssh/sshd_config file that include any setting other than VERBOSE or INFO:

```
# grep -Pi '^h*LogLevel\b' /etc/ssh/sshd_config
/etc/ssh/sshd_config.d/*.conf | grep -Evi '(VERBOSE|INFO)' | while read -r
l_out; do sed -ri "/^s*LogLevel\s+/s/^/# /" "$(awk -F: '{print $1}' <<<
$l_out)";done
```

Default Value:

LogLevel INFO

References:

1. https://www.ssh.com/ssh/sshd_config/
2. NIST SP 800-53 Rev. 5: AU-2, AU-12, SI-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 Collect Audit Logs Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.			
v7	6.2 Activate audit logging Ensure that local logging has been enabled on all systems and networking devices.			
v7	6.3 Enable Detailed Logging Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0005	

4.2.6 Ensure SSH PAM is enabled (Automated)

Profile Applicability:

- Level 1 - Server

Description:

UsePAM Enables the Pluggable Authentication Module interface. If set to “yes” this will enable PAM authentication using ChallengeResponseAuthentication and PasswordAuthentication in addition to PAM account and session module processing for all authentication types

Rationale:

When usePAM is set to yes, PAM runs through account and session types properly. This is important if you want to restrict access to services based off of IP, time or other factors of the account. Additionally, you can make sure users inherit certain environment variables on login or disallow access to the server

Impact:

If UsePAM is enabled, you will not be able to run sshd as a non-root user.

Audit:

Run the following command and verify that output matches:

```
# sshd -T -C user=root -C host="$(hostname)" -C addr="$(grep $(hostname) /etc/hosts | awk '{print $1}')" | grep -i usepam  
usepam yes
```

Run the following command and verify the output:

```
# grep -Pi '^h*UsePAM\b' /etc/ssh/sshd_config /etc/ssh/sshd_config.d/*.conf | grep -Evi 'yes'  
Nothing should be returned
```

Remediation:

Edit or create a file in the directory `/etc/ssh/sshd_config.d/` ending in `*.conf` or the `/etc/ssh/sshd_config` file and set the parameter as follows:

```
UsePAM yes
```

Run the following command to comment out any `UsePAM` parameter entries in files ending in `*.conf` in the `/etc/ssh/sshd_config.d/` directory or the `/etc/ssh/sshd_config` file that include any setting other than `yes`

```
# grep -Pi '^h*UsePAM\b' /etc/ssh/sshd_config /etc/ssh/sshd_config.d/*.conf  
| grep -Evi 'yes' | while read -r l_out; do sed -ri "/^s*UsePAM\s+/s/^/# /"  
"$ (awk -F: '{print $1}' <<< $l_out)";done
```

Default Value:

usePAM yes

References:

1. SSHD_CONFIG(5)
2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1021, T1021.004	TA0001	M1035

4.2.7 Ensure SSH root login is disabled (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `PermitRootLogin` parameter specifies if the root user can log in using ssh. The default is `no`.

Rationale:

Disallowing root logins over SSH requires system admins to authenticate using their own individual account, then escalating to root via `sudo` or `su`. This in turn limits opportunity for non-repudiation and provides a clear audit trail in the event of a security incident

Audit:

Run the following command and verify that output matches:

```
# sshd -T -C user=root -C host="$(hostname)" -C addr="$(grep $(hostname) /etc/hosts | awk '{print $1}')" | grep permitrootlogin
permitrootlogin no
```

Run the following command and verify the output:

```
# grep -Pi '^h*PermitRootLogin\b' /etc/ssh/sshd_config
/etc/ssh/sshd_config.d/*.conf | grep -Evi 'no'
Nothing should be returned
```

Remediation:

Edit or create a file ending in `*.conf` in the `/etc/ssh/sshd_config.d/` directory or the `/etc/ssh/sshd_config` file and set the `PermitRootLogin` parameter as follows:

```
PermitRootLogin no
```

Run the following command to comment out any `PermitRootLogin` parameter entries in files ending in `*.conf` in the `/etc/ssh/sshd_config.d/` directory or the `/etc/ssh/sshd_config` file that include any setting other than `no`

```
# grep -Pi '^h*PermitRootLogin\b' /etc/ssh/sshd_config
/etc/ssh/sshd_config.d/*.conf | grep -Evi 'no' | while read -r l_out; do sed
-ri "/^s*PermitRootLogin\s+/s/^/# /" "$(awk -F: '{print $1}' <<<
$l_out)";done
```

Default Value:

`PermitRootLogin without-password`

References:

1. SSHD_CONFIG(5)
2. NIST SP 800-53 Rev. 5: AC-6

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.4 <u>Restrict Administrator Privileges to Dedicated Administrator Accounts</u> Restrict administrator privileges to dedicated administrator accounts on enterprise assets. Conduct general computing activities, such as internet browsing, email, and productivity suite use, from the user's primary, non-privileged account.			
v7	4.3 <u>Ensure the Use of Dedicated Administrative Accounts</u> Ensure that all users with administrative account access use a dedicated or secondary account for elevated activities. This account should only be used for administrative activities and not internet browsing, email, or similar activities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1078	TA0001	M1026

4.2.8 Ensure SSH HostbasedAuthentication is disabled (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `HostbasedAuthentication` parameter specifies if authentication is allowed through trusted hosts via the user of `.rhosts`, or `/etc/hosts.equiv`, along with successful public key client host authentication. This option only applies to SSH Protocol Version 2.

Rationale:

Even though the `.rhosts` files are ineffective if support is disabled in `/etc/pam.conf`, disabling the ability to use `.rhosts` files in SSH provides an additional layer of protection.

Audit:

Run the following command and verify that output matches:

```
# sshd -T -C user=root -C host="$(hostname)" -C addr="$(grep $(hostname) /etc/hosts | awk '{print $1}')" | grep hostbasedauthentication  
hostbasedauthentication no
```

Run the following command and verify the output:

```
# grep -Pi '^h*HostbasedAuthentication\b' /etc/ssh/sshd_config  
/etc/ssh/sshd_config.d/*.conf | grep -Evi 'no'  
Nothing should be returned
```

Remediation:

Edit or create a file ending in *.conf in the /etc/ssh/sshd_config.d/ directory or the /etc/ssh/sshd_config file and set the HostbasedAuthentication parameter as follows:

```
HostbasedAuthentication no
```

Run the following command to comment out any HostbasedAuthentication parameter entries in files ending in *.conf in the /etc/ssh/sshd_config.d/ directory or the /etc/ssh/sshd_config file that include any setting other than no:

```
# grep -Pi '^\\h*HostbasedAuthentication\\b' /etc/ssh/sshd_config  
/etc/ssh/sshd_config.d/*.conf | grep -Evi 'no' | while read -r l_out; do sed  
-ri "/^\\s*HostbasedAuthentication\\s+\\s/^/# /" "$(awk -F: '{print $1}' <<<  
$l_out)";done
```

Default Value:

HostbasedAuthentication no

References:

1. SSHD_CONFIG(5)
2. NIST SP 800-53 Rev. 5: CM-1,CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.1 Establish and Maintain a Secure Configuration Process Establish and maintain a secure configuration process for enterprise assets (end-user devices, including portable and mobile, non-computing/IoT devices, and servers) and software (operating systems and applications). Review and update documentation annually, or when significant enterprise changes occur that could impact this Safeguard.			
v7	16.3 Require Multi-factor Authentication Require multi-factor authentication for all user accounts, on all systems, whether managed onsite or by a third-party provider.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1078, T1078.001, T1078.003	TA0008	M1042

4.2.9 Ensure SSH PermitEmptyPasswords is disabled (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `PermitEmptyPasswords` parameter specifies if the SSH server allows login to accounts with empty password strings.

Rationale:

Disallowing remote shell access to accounts that have an empty password reduces the probability of unauthorized access to the system

Audit:

Run the following command and verify that output matches:

```
# sshd -T -C user=root -C host="$(hostname)" -C addr="$(grep $(hostname) /etc/hosts | awk '{print $1}')" | grep permitemptypasswords  
permitemptypasswords no
```

Run the following command and verify the output:

```
# grep -Pi '^h*PermitEmptyPasswords\b' /etc/ssh/sshd_config  
/etc/ssh/sshd_config.d/*.conf | grep -Evi 'no'  
Nothing should be returned
```

Remediation:

Edit or create a file ending in `*.conf` in the `/etc/ssh/sshd_config.d/` directory or the `/etc/ssh/sshd_config` file and set the `PermitEmptyPasswords` parameter as follows:

```
PermitEmptyPasswords no
```

Run the following command to comment out any `PermitEmptyPasswords` parameter entries in files ending in `*.conf` in the `/etc/ssh/sshd_config.d/` directory or the `/etc/ssh/sshd_config` file that include any setting other than `no`

```
# grep -Pi '^h*PermitEmptyPasswords\b' /etc/ssh/sshd_config  
/etc/ssh/sshd_config.d/*.conf | grep -Evi 'no' | while read -r l_out; do sed  
-ri "/^s*PermitEmptyPasswords\s+/s/^/# /" "$l_out" <<<  
$(awk -F: '{print $1}' <<<  
$l_out)";done
```

Default Value:

`PermitEmptyPasswords no`

References:

1. SSHD_CONFIG(5)
2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.2 Use Unique Passwords Use unique passwords for all enterprise assets. Best practice implementation includes, at a minimum, an 8-character password for accounts using MFA and a 14-character password for accounts not using MFA.			
v7	4.4 Use Unique Passwords Where multi-factor authentication is not supported (such as local administrator, root, or service accounts), accounts will use passwords that are unique to that system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1021, T1021.000	TA0008	M1042

4.2.10 Ensure SSH PermitUserEnvironment is disabled (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `PermitUserEnvironment` option allows users to present environment options to the `ssh` daemon.

Rationale:

Permitting users the ability to set environment variables through the SSH daemon could potentially allow users to bypass security controls (e.g. setting an execution path that has `ssh` executing trojan'd programs)

Audit:

Run the following command and verify that output matches:

```
# sshd -T -C user=root -C host="$(hostname)" -C addr="$(grep $(hostname) /etc/hosts | awk '{print $1}')" | grep permituserenvironment  
permituserenvironment no
```

Run the following command and verify the output:

```
# grep -Pi '^h*PermitUserEnvironment\b' /etc/ssh/sshd_config  
/etc/ssh/sshd_config.d/*.conf | grep -Evi 'no'  
Nothing should be returned
```

Remediation:

Edit or create a file ending in `*.conf` in the `/etc/ssh/sshd_config.d/` directory or the `/etc/ssh/sshd_config` file and set the `PermitUserEnvironment` parameter as follows:

```
PermitUserEnvironment no
```

Run the following command to comment out any `PermitUserEnvironment` parameter entries in files ending in `*.conf` in the `/etc/ssh/sshd_config.d/` directory or the `/etc/ssh/sshd_config` file that include any setting other than `no`

```
# grep -Pi '^h*PermitUserEnvironment\b' /etc/ssh/sshd_config  
/etc/ssh/sshd_config.d/*.conf | grep -Evi 'no' | while read -r l_out; do sed  
-ri "/^s*PermitUserEnvironment\s+/s/^/# /" "$(awk -F: '{print $1}' <<<  
$l_out)";done
```

Default Value:

`PermitUserEnvironment no`

References:

1. SSHD_CONFIG(5)
2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.1 <u>Establish and Maintain a Secure Configuration Process</u> Establish and maintain a secure configuration process for enterprise assets (end-user devices, including portable and mobile, non-computing/IoT devices, and servers) and software (operating systems and applications). Review and update documentation annually, or when significant enterprise changes occur that could impact this Safeguard.			
v7	5.1 <u>Establish Secure Configurations</u> Maintain documented, standard security configuration standards for all authorized operating systems and software.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1021, T1021.000	TA0008	M1042

4.2.11 Ensure SSH IgnoreRhosts is enabled (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `IgnoreRhosts` parameter specifies that `.rhosts` and `.shosts` files will not be used in `RhostsRSAAuthentication` OR `HostbasedAuthentication`.

Rationale:

Setting this parameter forces users to enter a password when authenticating with `ssh`.

Audit:

Run the following command and verify that output matches:

```
# sshd -T -C user=root -C host="$(hostname)" -C addr="$(grep $(hostname) /etc/hosts | awk '{print $1}')" | grep ignorerhosts  
ignorerhosts yes
```

Run the following command and verify the output:

```
# grep -Pi '^h*ignorerhosts\b' /etc/ssh/sshd_config  
/etc/ssh/sshd_config.d/*.conf | grep -Evi 'yes'  
  
Nothing should be returned
```

Remediation:

Edit or create a file ending in `*.conf` in the `/etc/ssh/sshd_config.d/` directory or the `/etc/ssh/sshd_config` file and set the `IgnoreRhosts` parameter as follows:

```
IgnoreRhosts yes
```

Run the following command to comment out any `IgnoreRhosts` parameter entries in files ending in `*.conf` in the `/etc/ssh/sshd_config.d/` directory or the `/etc/ssh/sshd_config` file that include any setting other than `yes`

```
# grep -Pi '^h*IgnoreRhosts\b' /etc/ssh/sshd_config  
/etc/ssh/sshd_config.d/*.conf | grep -Evi 'yes' | while read -r l_out; do sed  
-ri "/^s*IgnoreRhosts\s+\/s\/^\/# /" "$l_out" <<< "$l_out";done
```

Default Value:

`IgnoreRhosts yes`

References:

1. SSHD_CONFIG(5)
2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1078, T1078.001, T1078.003	TA0001	M1027

4.2.12 Ensure SSH X11 forwarding is disabled (Automated)

Profile Applicability:

- Level 2 - Server

Description:

The `X11Forwarding` parameter provides the ability to tunnel X11 traffic through the connection to enable remote graphic connections.

Rationale:

Disable X11 forwarding unless there is an operational requirement to use X11 applications directly. There is a small risk that the remote X11 servers of users who are logged in via SSH with X11 forwarding could be compromised by other users on the X11 server. Note that even if X11 forwarding is disabled, users can always install their own forwarders.

Audit:

Run the following command and verify that output matches:

```
# sshd -T -C user=root -C host="$(hostname)" -C addr="$(grep $(hostname) /etc/hosts | awk '{print $1}')" | grep -i x11forwarding
x11forwarding no
```

Run the following command and verify the output:

```
# grep -Pi '^h*X11Forwarding\b' /etc/ssh/sshd_config
/etc/ssh/sshd_config.d/*.conf | grep -Evi 'no'
Nothing should be returned
```

Remediation:

Edit or create a file ending in `*.conf` in the `/etc/ssh/sshd_config.d/` directory or the `/etc/ssh/sshd_config` file and set the `X11Forwarding` parameter as follows:

```
X11Forwarding no
```

Run the following command to comment out any `X11Forwarding` parameter entries in files ending in `*.conf` in the `/etc/ssh/sshd_config.d/` directory or the `/etc/ssh/sshd_config` file that include any setting other than `no`

```
# grep -Pi '^h*X11Forwarding\b' /etc/ssh/sshd_config
/etc/ssh/sshd_config.d/*.conf | grep -Evi 'no' | while read -r l_out; do sed
-ri "/^s*X11Forwarding\s+/s/^/# /" "$(awk -F: '{print $1}' <<< $l_out)";done
```

References:

1. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v7	<u>9.2 Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1210, T1210.000	TA0008	M1042

4.2.13 Ensure SSH AllowTcpForwarding is disabled (Automated)

Profile Applicability:

- Level 2 - Server

Description:

SSH port forwarding is a mechanism in SSH for tunneling application ports from the client to the server, or servers to clients. It can be used for adding encryption to legacy applications, going through firewalls, and some system administrators and IT professionals use it for opening backdoors into the internal network from their home machines

Rationale:

Leaving port forwarding enabled can expose the organization to security risks and backdoors.

SSH connections are protected with strong encryption. This makes their contents invisible to most deployed network monitoring and traffic filtering solutions. This invisibility carries considerable risk potential if it is used for malicious purposes such as data exfiltration. Cybercriminals or malware could exploit SSH to hide their unauthorized communications, or to exfiltrate stolen data from the target network

Impact:

SSH tunnels are widely used in many corporate environments that employ mainframe systems as their application backends. In those environments the applications themselves may have very limited native support for security. By utilizing tunneling, compliance with SOX, HIPAA, PCI-DSS, and other standards can be achieved without having to modify the applications.

Audit:

Run the following command and verify that output matches:

```
# sshd -T -C user=root -C host="$(hostname)" -C addr="$(grep $(hostname) /etc/hosts | awk '{print $1}')" | grep -i allowtcpforwarding  
  
allowtcpforwarding no
```

Run the following command and verify the output:

```
# grep -Pi '^h*AllowTcpForwarding\b' /etc/ssh/sshd_config  
/etc/ssh/sshd_config.d/*.conf | grep -Evi 'no'  
  
Nothing should be returned
```

Remediation:

Edit or create a file ending in *.conf in the /etc/ssh/sshd_config.d/ directory or the /etc/ssh/sshd_config file and set the AllowTcpForwarding parameter as follows:

```
AllowTcpForwarding no
```

Run the following command to comment out any AllowTcpForwarding parameter entries in files ending in *.conf in the /etc/ssh/sshd_config.d/ directory or the /etc/ssh/sshd_config file that include any setting other than no:

```
# grep -Pi '^h*AllowTcpForwarding\b' /etc/ssh/sshd_config  
/etc/ssh/sshd_config.d/*.conf | grep -Evi 'no' | while read -r l_out; do sed  
-ri "/^s*AllowTcpForwarding\s+/s/^/# /" "$(awk -F: '{print $1}' <<<  
$l_out)";done
```

Default Value:

AllowTcpForwarding yes

References:

1. <https://www.ssh.com/ssh/tunneling/example>
2. NIST SP 800-53 Rev. 5: CM-1,CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.1 <u>Establish and Maintain a Secure Configuration Process</u> Establish and maintain a secure configuration process for enterprise assets (end-user devices, including portable and mobile, non-computing/IoT devices, and servers) and software (operating systems and applications). Review and update documentation annually, or when significant enterprise changes occur that could impact this Safeguard.			
v7	9.2 <u>Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1048, T1048.002, T1572, T1572.000	TA0008	M1042

4.2.14 Ensure system-wide crypto policy is not over-ridden (Automated)

Profile Applicability:

- Level 1 - Server

Description:

System-wide Crypto policy can be over-ridden or opted out of for openSSH

Rationale:

Over-riding or opting out of the system-wide crypto policy could allow for the use of less secure Ciphers, MACs, KexAlgorithms and GSSAPIKexAlgorithm

Audit:

Run the following command:

```
# grep -i '^s*CRYPTO_POLICY=' /etc/sysconfig/ssh  
/etc/ssh/ssh_config.d/*.conf
```

No output should be returned

Remediation:

Run the following commands:

```
# sed -ri "s/^(CRYPTO_POLICY\s*=.*)$/# \1/" /etc/sysconfig/ssh  
/etc/ssh/ssh_config.d/*.conf  
  
# systemctl reload sshd
```

References:

1. NIST SP 800-53 Rev. 5: SC-8, IA-5, AC-17

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.10 <u>Encrypt Sensitive Data in Transit</u> Encrypt sensitive data in transit. Example implementations can include: Transport Layer Security (TLS) and Open Secure Shell (OpenSSH).		●	●
v7	14.4 <u>Encrypt All Sensitive Information in Transit</u> Encrypt all sensitive information in transit.		●	●

4.2.15 Ensure SSH warning banner is configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `Banner` parameter specifies a file whose contents must be sent to the remote user before authentication is permitted. By default, no banner is displayed.

Rationale:

Banners are used to warn connecting users of the particular site's policy regarding connection. Presenting a warning message prior to the normal user login may assist the prosecution of trespassers on the computer system.

Audit:

Run the following command and verify that output matches:

```
# sshd -T -C user=root -C host="$(hostname)" -C addr="$(grep $(hostname) /etc/hosts | awk '{print $1}')" | grep banner  
  
banner /etc/issue.net
```

Remediation:

Edit or create a file ending in `*.conf` in the `/etc/ssh/sshd_config.d/` directory or the `/etc/ssh/sshd_config` file and set the `Banner` parameter as follows:

```
Banner /etc/issue.net
```

References:

1. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1082, T1082.000, T1592, T1592.004	TA0001, TA0007	M1035

4.2.16 Ensure SSH MaxAuthTries is set to 4 or less (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `MaxAuthTries` parameter specifies the maximum number of authentication attempts permitted per connection. When the login failure count reaches half the number, error messages will be written to the `syslog` file detailing the login failure.

Rationale:

Setting the `MaxAuthTries` parameter to a low number will minimize the risk of successful brute force attacks to the SSH server. While the recommended setting is 4, set the number based on site policy.

Audit:

Run the following command and verify that output `MaxAuthTries` is 4 or less:

```
# sshd -T -C user=root -C host="$(hostname)" -C addr="$(grep $(hostname) /etc/hosts | awk '{print $1}')" | grep maxauthtries
maxauthtries 4
```

Run the following command and verify that the output:

```
# grep -Pi '^h*maxauthtries\h+([5-9]|[1-9][0-9]+)' /etc/ssh/sshd_config /etc/ssh/sshd_config.d/*.conf
Nothing is returned
```

Remediation:

Edit or create a file ending in `*.conf` in the `/etc/ssh/sshd_config.d/` directory or the `/etc/ssh/sshd_config` file and set the `MaxAuthTries` parameter as follows:

```
MaxAuthTries 4
```

Run the following command to comment out any `MaxAuthTries` parameter entries in files ending in `*.conf` in the `/etc/ssh/sshd_config.d/` directory or the `/etc/ssh/sshd_config` file that include any setting greater than 4:

```
# grep -Pi '^h*maxauthtries\h+([5-9]|[1-9][0-9]+)' /etc/ssh/sshd_config /etc/ssh/sshd_config.d/*.conf | while read -r l_out; do sed -ri "/^s*maxauthtries\s+([5-9]|[1-9][0-9]+)/s/^/# /" "$l_out" <<< $(awk -F: '{print $1}' <<< $l_out);done
```

Default Value:

`MaxAuthTries 6`

References:

1. SSHD_CONFIG(5)
2. NIST SP 800-53 Rev. 5: AU-3

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 <u>Collect Detailed Audit Logs</u> Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.		●	●
v7	16.13 <u>Alert on Account Login Behavior Deviation</u> Alert when users deviate from normal login behavior, such as time-of-day, workstation location and duration.			●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1110, T1110.001, T1110.003	TA0006	M1036

4.2.17 Ensure SSH MaxStartups is configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `MaxStartups` parameter specifies the maximum number of concurrent unauthenticated connections to the SSH daemon.

Rationale:

To protect a system from denial of service due to a large number of pending authentication connection attempts, use the rate limiting function of `MaxStartups` to protect availability of `sshd` logins and prevent overwhelming the daemon.

Audit:

Run the following command and verify that output `MaxStartups` is `10:30:60` or more restrictive:

```
# sshd -T -C user=root -C host="$(hostname)" -C addr="$(grep $(hostname) /etc/hosts | awk '{print $1}')" | grep -i maxstartups
maxstartups 10:30:60
```

Run the following command and verify the output:

```
# grep -Ei '^\s*maxstartups\s+(((1[0-9]|[0-9][0-9]+):([0-9]+):([0-9]+))|((([0-9]+):(3[0-9]|[4-9][0-9]|[1-9][0-9][0-9]+):([0-9]+))|((([0-9]+):([0-9]+):(6[0-9]|[7-9][0-9]|[1-9][0-9][0-9]+)))' /etc/ssh/sshd_config
/etc/ssh/sshd_config.d/*.conf
Nothing should be returned
```

Remediation:

Edit or create a file ending in *.conf in the /etc/ssh/sshd_config.d/ directory or the /etc/ssh/sshd_config file and set the MaxStartups parameter as follows:

```
MaxStartups 10:30:60
```

Run the following command to comment out any MaxStartups parameter entries in files ending in *.conf in the /etc/ssh/sshd_config.d/ directory or the /etc/ssh/sshd_config file that include any setting greater than 10:30:60:

```
# grep -Pi '\s*maxstartups\s+(((1[0-9]|[0-9][0-9]+):([0-9]+):([0-9]+))|((([0-9]+):(3[0-9]|[4-9][0-9]|[1-9][0-9][0-9]+):([0-9]+))|((([0-9]+):([0-9]+):(6[0-9]|[7-9][0-9]|[1-9][0-9][0-9]+))))' /etc/ssh/sshd_config
/etc/ssh/sshd_config.d/*.conf | while read -r l_out; do sed -ri
"/^\s*MaxStartups\s+(((1[0-9]|[0-9][0-9]+):([0-9]+):([0-9]+))|((([0-9]+):(3[0-9]|[4-9][0-9]|[1-9][0-9][0-9]+):([0-9]+))|((([0-9]+):([0-9]+):(6[0-9]|[7-9][0-9]|[1-9][0-9][0-9]+))))/s/^/# /" "$(awk -F: '{print $1}' <<<
$l_out)";done
```

Default Value:

MaxStartups 10:30:100

References:

1. SSHD_CONFIG(5)
2. NIST SP 800-53 Rev. 5: CM-1,CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.1 Establish and Maintain a Secure Configuration Process Establish and maintain a secure configuration process for enterprise assets (end-user devices, including portable and mobile, non-computing/IoT devices, and servers) and software (operating systems and applications). Review and update documentation annually, or when significant enterprise changes occur that could impact this Safeguard.			
v7	5.1 Establish Secure Configurations Maintain documented, standard security configuration standards for all authorized operating systems and software.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1499, T1499.002	TA0040	M1036

4.2.18 Ensure SSH MaxSessions is set to 10 or less (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `MaxSessions` parameter specifies the maximum number of open sessions permitted from a given connection.

Rationale:

To protect a system from denial of service due to a large number of concurrent sessions, use the rate limiting function of `MaxSessions` to protect availability of `ssh` logins and prevent overwhelming the daemon.

Audit:

Run the following command and verify that output `MaxSessions` is 10 or less:

```
# sshd -T -C user=root -C host="$(hostname)" -C addr="$(grep $(hostname)
/etc/hosts | awk '{print $1}')" | grep -i maxsessions

maxsessions 10
```

Run the following command and verify the output:

```
grep -Ei '^s*MaxSessions\s+(1[1-9]|[2-9][0-9]|[1-9][0-9][0-9]+)'
/etc/ssh/sshd_config /etc/ssh/sshd_config.d/*.conf

Nothing should be returned
```

Remediation:

Edit or create a file ending in `*.conf` in the `/etc/ssh/sshd_config.d/` directory or the `/etc/ssh/sshd_config` file and set the `MaxSessions` parameter as follows:

```
MaxSessions 10
```

Run the following command to comment out any `MaxSessions` parameter entries in files ending in `*.conf` in the `/etc/ssh/sshd_config.d/` directory or the `/etc/ssh/sshd_config` file that include any setting greater than 10

```
# grep -Pi '^s*MaxSessions\s+(1[1-9]|[2-9][0-9]|[1-9][0-9][0-9]+)'
/etc/ssh/sshd_config /etc/ssh/sshd_config.d/*.conf | while read -r l_out; do
sed -ri "/^s*MaxSessions\s+(1[1-9]|[2-9][0-9]|[1-9][0-9][0-9]+)/s/^/# /"
"${awk -F: '{print $1}' <<< $l_out}";done
```

Default Value:

`MaxSessions 10`

References:

1. SSHD_CONFIG(5)
2. NIST SP 800-53 Rev. 5: CM-1,CM-2, CM-6, CM-7, IA-5

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1499, T1499.002	TA0040	M1036

4.2.19 Ensure SSH LoginGraceTime is set to one minute or less (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `LoginGraceTime` parameter specifies the time allowed for successful authentication to the SSH server. The longer the Grace period is the more open unauthenticated connections can exist. Like other session controls in this session the Grace Period should be limited to appropriate organizational limits to ensure the service is available for needed access.

Rationale:

Setting the `LoginGraceTime` parameter to a low number will minimize the risk of successful brute force attacks to the SSH server. It will also limit the number of concurrent unauthenticated connections. While the recommended setting is 60 seconds (1 Minute), set the number based on site policy.

Audit:

Run the following command and verify that output `LoginGraceTime` is between 1 and 60 seconds or 1m:

```
# sshd -T -C user=root -C host="$(hostname)" -C addr="$(grep $(hostname) /etc/hosts | awk '{print $1}')" | grep logingracetime  
logingracetime 60
```

Run the following command and verify the output:

```
# grep -Ei '^s*LoginGraceTime\s+(0|6[1-9]|[7-9][0-9]|[1-9][0-9][0-9]+|^[^1]m)' /etc/ssh/sshd_config /etc/ssh/sshd_config.d/*.conf  
Nothing should be returned
```

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the parameter as follows:

```
LoginGraceTime 60
```

Edit or create a file ending in `*.conf` in the `/etc/ssh/sshd_config.d/` directory or the `/etc/ssh/sshd_config` file and set the `LoginGraceTime` parameter as follows:

```
LoginGraceTime 60  
-or-  
LoginGraceTime 1m
```

Run the following command to comment out any `LoginGraceTime` parameter entries in files ending in `*.conf` in the `/etc/ssh/sshd_config.d/` directory or the `/etc/ssh/sshd_config` file that include any setting equal to 0 or greater than 60 seconds:

```
# grep -Pi '^s*LoginGraceTime\s+(0|6[1-9]|[7-9][0-9]|[1-9][0-9][0-9]+|^[^1]m)' /etc/ssh/sshd_config /etc/ssh/sshd_config.d/*.conf | while read -r l_out; do sed -ri "/^s*LoginGraceTime\s+(0|6[1-9]|[7-9][0-9]|[1-9][0-9][0-9]+|^[^1]m)/s/^/# /" "$(awk -F: '{print $1}' <<< $l_out)";done
```

Default Value:

`LoginGraceTime 120`

References:

1. SSHD_CONFIG(5)
2. NIST SP 800-53 Rev. 5: CM-6

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1110, T1110.001, T1110.003, T1110.004	TA0006	M1036

4.2.20 Ensure SSH Idle Timeout Interval is configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

NOTE: To clarify, the two settings described below are only meant for idle connections from a protocol perspective and not meant to check if the user is active or not. An idle user does not mean an idle connection. SSH does not, and never had, intentionally the capability to drop idle users. In SSH versions before 8.2p1 there was a bug that caused these values to behave in such a manner that they were abused to disconnect idle users. This bug has been resolved in 8.2p1 and thus may no longer be abused to disconnect idle users.

The two options `ClientAliveInterval` and `ClientAliveCountMax` control the timeout of SSH sessions. Taken directly from `man 5 sshd_config`:

- `ClientAliveInterval` Sets a timeout interval in seconds after which if no data has been received from the client, `sshd(8)` will send a message through the encrypted channel to request a response from the client. The default is 0, indicating that these messages will not be sent to the client.
- `ClientAliveCountMax` Sets the number of client alive messages which may be sent without `sshd(8)` receiving any messages back from the client. If this threshold is reached while client alive messages are being sent, `sshd` will disconnect the client, terminating the session. It is important to note that the use of client alive messages is very different from `TCPKeepAlive`. The client alive messages are sent through the encrypted channel and therefore will not be spoofable. The TCP keepalive option enabled by `TCPKeepAlive` is spoofable. The client alive mechanism is valuable when the client or server depend on knowing when a connection has become unresponsive. The default value is 3. If `ClientAliveInterval` is set to 15, and `ClientAliveCountMax` is left at the default, unresponsive SSH clients will be disconnected after approximately 45 seconds. Setting a zero `ClientAliveCountMax` disables connection termination.

Rationale:

In order to prevent resource exhaustion, appropriate values should be set for both `ClientAliveInterval` and `ClientAliveCountMax`. Specifically, looking at the source code, `ClientAliveCountMax` must be greater than zero in order to utilize the ability of SSH to drop idle connections. If connections are allowed to stay open indefinitely, this can potentially be used as a DDOS attack or simple resource exhaustion could occur over unreliable networks.

The example set here is a 45 second timeout. Consult your site policy for network timeouts and apply as appropriate.

Audit:

Run the following commands and verify `ClientAliveInterval` is greater than zero:

```
# sshd -T -C user=root -C host="$(hostname)" -C addr="$(grep $(hostname) /etc/hosts | awk '{print $1}')" | grep clientaliveinterval
```

Example output:

```
clientaliveinterval 15
```

Run the following command and verify `ClientAliveCountMax` is greater than zero:

```
# sshd -T -C user=root -C host="$(hostname)" -C addr="$(grep $(hostname) /etc/hosts | awk '{print $1}')" | grep clientalivecountmax
```

Example output:

```
clientalivecountmax 3
```

Remediation:

Edit or create a file ending in `*.conf` in the `/etc/ssh/sshd_config.d/` directory or the `/etc/ssh/sshd_config` file and set the `ClientAliveInterval` and `ClientAliveCountMax` parameters according to site policy.

Example:

```
ClientAliveInterval 15
ClientAliveCountMax 3
```

Edit files ending in `*.conf` in the `/etc/ssh/sshd_config.d/` directory and the `/etc/ssh/sshd_config` file and remove occurrences of the `ClientAliveInterval` and `ClientAliveCountMax` parameters not in accordance with local site policy.

Run the following command to comment out any `ClientAliveCountMax` parameter entries in files ending in `*.conf` in the `/etc/ssh/sshd_config.d/` directory or the `/etc/ssh/sshd_config` file that include the setting of 0 "disabled":

```
# grep -Pi '^h*ClientAliveCountMax\h+0\b' /etc/ssh/sshd_config
/etc/ssh/sshd_config.d/*.conf | while read -r l_out; do sed -ri
"/^s*ClientAliveCountMax\s+0/s/^/# /" "$(awk -F: '{print $1}' <<<
$l_out)";done
```

Default Value:

`ClientAliveInterval` 0

`ClientAliveCountMax` 3

References:

1. https://man.openbsd.org/sshd_config
2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

Additional Information:

https://bugzilla.redhat.com/show_bug.cgi?id=1873547

https://github.com/openssh/openssh-portable/blob/V_8_9/serverloop.c#L137

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1078, T1078.001, T1078.002, T1078.003	TA0001	M1026

4.3 Configure privilege escalation

There are various tools which allows a permitted user to execute a command as the superuser or another user, as specified by the security policy.

sudo

<https://www.sudo.ws/>

The invoking user's real (not effective) user ID is used to determine the user name with which to query the security policy.

`sudo` supports a plug-in architecture for security policies and input/output logging. Third parties can develop and distribute their own policy and I/O logging plug-ins to work seamlessly with the `sudo` front end. The default security policy is `sudoers`, which is configured via the file `/etc/sudoers` and any entries in `/etc/sudoers.d`.

pkexec

<https://www.freedesktop.org/software/polkit/docs/0.105/pkexec.1.html>

4.3.1 Ensure sudo is installed (Automated)

Profile Applicability:

- Level 1 - Server

Description:

`sudo` allows a permitted user to execute a command as the superuser or another user, as specified by the security policy. The invoking user's real (not effective) user ID is used to determine the user name with which to query the security policy.

Rationale:

`sudo` supports a plug-in architecture for security policies and input/output logging. Third parties can develop and distribute their own policy and I/O logging plug-ins to work seamlessly with the `sudo` front end. The default security policy is `sudoers`, which is configured via the file `/etc/sudoers` and any entries in `/etc/sudoers.d`.

The security policy determines what privileges, if any, a user has to run `sudo`. The policy may require that users authenticate themselves with a password or another authentication mechanism. If authentication is required, `sudo` will exit if the user's password is not entered within a configurable time limit. This limit is policy-specific.

Audit:

Verify that `sudo` is installed.

Run the following command:

```
# dnf list sudo

Installed Packages
sudo.x86_64          <VERSION>          @anaconda
Available Packages
sudo.x86_64          <VERSION>          updates
```

Remediation:

Run the following command to install `sudo`

```
# dnf install sudo
```

References:

1. SUDO(8)
2. NIST SP 800-53 Rev. 5: AC-6

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.4 <u>Restrict Administrator Privileges to Dedicated Administrator Accounts</u> Restrict administrator privileges to dedicated administrator accounts on enterprise assets. Conduct general computing activities, such as internet browsing, email, and productivity suite use, from the user's primary, non-privileged account.			
v7	4.3 <u>Ensure the Use of Dedicated Administrative Accounts</u> Ensure that all users with administrative account access use a dedicated or secondary account for elevated activities. This account should only be used for administrative activities and not internet browsing, email, or similar activities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1548	TA0004	M1026

4.3.2 Ensure sudo commands use pty (Automated)

Profile Applicability:

- Level 1 - Server

Description:

`sudo` can be configured to run only from a pseudo terminal (`pseudo-pty`).

Rationale:

Attackers can run a malicious program using `sudo` which would fork a background process that remains even when the main program has finished executing.

Impact:

WARNING: Editing the `sudo` configuration incorrectly can cause `sudo` to stop functioning. Always use `visudo` to modify `sudo` configuration files.

Audit:

Verify that `sudo` can only run other commands from a pseudo terminal.

Run the following command:

```
# grep -rPi '^h*Defaults\h+([\^#\n\r]+,)?use_pty(,\h*\h+\h*)*\h*(#.*)?$'
/etc/sudoers*

/etc/sudoers:Defaults use_pty
```

Remediation:

Edit the file `/etc/sudoers` with `visudo` or a file in `/etc/sudoers.d/` with `visudo -f <PATH_TO_FILE>` and add the following line:

```
Defaults use_pty
```

Note:

- `sudo` will read each file in `/etc/sudoers.d`, skipping file names that end in `~` or contain a `.` character to avoid causing problems with package manager or editor temporary/backup files.
- Files are parsed in sorted lexical order. That is, `/etc/sudoers.d/01_first` will be parsed before `/etc/sudoers.d/10_second`.
- Be aware that because the sorting is lexical, not numeric, `/etc/sudoers.d/1_whoops` would be loaded after `/etc/sudoers.d/10_second`.
- Using a consistent number of leading zeroes in the file names can be used to avoid such problems.

References:

1. SUDO(8)
2. VISUDO(8)

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.4 <u>Restrict Administrator Privileges to Dedicated Administrator Accounts</u> Restrict administrator privileges to dedicated administrator accounts on enterprise assets. Conduct general computing activities, such as internet browsing, email, and productivity suite use, from the user's primary, non-privileged account.			
v7	5.1 <u>Establish Secure Configurations</u> Maintain documented, standard security configuration standards for all authorized operating systems and software.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1078, T1078.003, T1548, T1548.003	TA0001, TA0003	M1026, M1038

4.3.3 Ensure sudo log file exists (Automated)

Profile Applicability:

- Level 1 - Server

Description:

sudo can use a custom log file

Rationale:

A sudo log file simplifies auditing of sudo commands

Impact:

WARNING: Editing the `sudo` configuration incorrectly can cause `sudo` to stop functioning. Always use `visudo` to modify `sudo` configuration files.

Creation of additional log files can cause disk space exhaustion if not correctly managed. You should configure `logrotate` to manage the sudo log in accordance with your local policy.

Audit:

Run the following command to verify that sudo has a custom log file configured

```
# grep -rPsi
"^\h*Defaults\h+([\^#]+,\h*)?logfile\h*=\h*(\"|\')?\H+(\"|\')?(,\h*\H+\h*)*\h*
(#.*)?$" /etc/sudoers*

Defaults logfile="/var/log/sudo.log"
```

Note:

- `sudo` will read each file in `/etc/sudoers.d`, skipping file names that end in `~` or contain a `.` character to avoid causing problems with package manager or editor temporary/backup files.
- Files are parsed in sorted lexical order. That is, `/etc/sudoers.d/01_first` will be parsed before `/etc/sudoers.d/10_second`.
- Be aware that because the sorting is lexical, not numeric, `/etc/sudoers.d/1_whoops` would be loaded after `/etc/sudoers.d/10_second`.
- Using a consistent number of leading zeroes in the file names can be used to avoid such problems.

Remediation:

Edit the file `/etc/sudoers` or a file in `/etc/sudoers.d/` with `visudo` or `visudo -f <PATH TO FILE>` and add the following line:

```
Defaults logfile="<PATH TO CUSTOM LOG FILE>"
```

Example

```
Defaults logfile="/var/log/sudo.log"
```

References:

1. SUDO(8)
2. VISUDO(8)
3. sudoers(5)

Additional Information:

`visudo` edits the `sudoers` file in a safe fashion, analogous to `vipw(8)`. `visudo` locks the `sudoers` file against multiple simultaneous edits, provides basic sanity checks, and checks for parse errors. If the `sudoers` file is currently being edited you will receive a message to try again later.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 Collect Detailed Audit Logs Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.			
v7	6.3 Enable Detailed Logging Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0004	M1026

4.3.4 Ensure re-authentication for privilege escalation is not disabled globally (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The operating system must be configured so that users must re-authenticate for privilege escalation.

Rationale:

Without re-authentication, users may access resources or perform tasks for which they do not have authorization.

When operating systems provide the capability to escalate a functional capability, it is critical the user re-authenticate.

Audit:

Verify the operating system requires users to re-authenticate for privilege escalation. Check the configuration of the `/etc/sudoers` and `/etc/sudoers.d/*` files with the following command:

```
# grep -r "^[^#].*\!authenticate" /etc/sudoers*
```

If any line is found with a `!authenticate` tag, refer to the remediation procedure below.

Remediation:

Configure the operating system to require users to reauthenticate for privilege escalation.

Based on the outcome of the audit procedure, use `visudo -f <PATH TO FILE>` to edit the relevant sudoers file.

Remove any occurrences of `!authenticate` tags in the file(s).

References:

1. NIST SP 800-53 Rev. 5: AC-6

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>5.4 Restrict Administrator Privileges to Dedicated Administrator Accounts</u> Restrict administrator privileges to dedicated administrator accounts on enterprise assets. Conduct general computing activities, such as internet browsing, email, and productivity suite use, from the user's primary, non-privileged account.			
v7	<u>4.3 Ensure the Use of Dedicated Administrative Accounts</u> Ensure that all users with administrative account access use a dedicated or secondary account for elevated activities. This account should only be used for administrative activities and not internet browsing, email, or similar activities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1548	TA0004	M1026

4.3.5 Ensure sudo authentication timeout is configured correctly (Automated)

Profile Applicability:

- Level 1 - Server

Description:

`sudo` caches used credentials for a default of 5 minutes. This is for ease of use when there are multiple administrative tasks to perform. The timeout can be modified to suit local security policies.

Rationale:

Setting a timeout value reduces the window of opportunity for unauthorized privileged access to another user.

Audit:

Ensure that the caching timeout is no more than 15 minutes.

Example:

```
# grep -roP "timestamp_timeout=\K[0-9]*" /etc/sudoers*
```

If there is no `timestamp_timeout` configured in `/etc/sudoers*` then the default is 5 minutes. This default can be checked with:

```
# sudo -V | grep "Authentication timestamp timeout:"
```

NOTE: A value of `-1` means that the timeout is disabled. Depending on the configuration of the `timestamp_type`, this could mean for all terminals / processes of that user and not just that one single terminal session.

Remediation:

If the currently configured timeout is larger than 15 minutes, edit the file listed in the audit section with `visudo -f <PATH TO FILE>` and modify the entry `timestamp_timeout=` to 15 minutes or less as per your site policy. The value is in minutes. This particular entry may appear on its own, or on the same line as `env_reset`. See the following two examples:

```
Defaults    env_reset, timestamp_timeout=15
Defaults    timestamp_timeout=15
Defaults    env_reset
```

References:

1. <https://www.sudo.ws/man/1.9.0/sudoers.man.html>
2. NIST SP 800-53 Rev. 5: AC-6

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>5.4 Restrict Administrator Privileges to Dedicated Administrator Accounts</u> Restrict administrator privileges to dedicated administrator accounts on enterprise assets. Conduct general computing activities, such as internet browsing, email, and productivity suite use, from the user's primary, non-privileged account.			
v7	<u>4.3 Ensure the Use of Dedicated Administrative Accounts</u> Ensure that all users with administrative account access use a dedicated or secondary account for elevated activities. This account should only be used for administrative activities and not internet browsing, email, or similar activities.			

4.3.6 Ensure access to the su command is restricted (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `su` command allows a user to run a command or shell as another user. The program has been superseded by `sudo`, which allows for more granular control over privileged access. Normally, the `su` command can be executed by any user. By uncommenting the `pam_wheel.so` statement in `/etc/pam.d/su`, the `su` command will only allow users in a specific groups to execute `su`. This group should be empty to reinforce the use of `sudo` for privileged access.

Rationale:

Restricting the use of `su`, and using `sudo` in its place, provides system administrators better control of the escalation of user privileges to execute privileged commands. The `sudo` utility also provides a better logging and audit mechanism, as it can log each command executed via `sudo`, whereas `su` can only record that a user executed the `su` program.

Audit:

Run the following command and verify the output matches the line:

```
# grep -Pi
'^\h*auth\h+(?:required|requisite)\h+pam_wheel\.so\h+(?:[^\n\r]+\h+)?((?!\\2)
(use_uid\b|group=\H+\b))\h+(?:[^\n\r]+\h+)?((?!\\1)(use_uid\b|group=\H+\b))(\
h+\.*)?$', /etc/pam.d/su

auth required pam_wheel.so use_uid group=<group_name>
```

Run the following command and verify that the group specified in `<group_name>` contains no users:

```
# grep <group_name> /etc/group

<group_name>:x:<GID>:
```

There should be no users listed after the Group ID field.

Remediation:

Create an empty group that will be specified for use of the `su` command. The group should be named according to site policy.

Example:

```
# groupadd sugroup
```

Add the following line to the `/etc/pam.d/su` file, specifying the empty group:

```
auth required pam_wheel.so use_uid group=sugroup
```

References:

1. NIST SP 800-53 Rev. 5: AC-3. MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	5.1 <u>Establish Secure Configurations</u> Maintain documented, standard security configuration standards for all authorized operating systems and software.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1078	TA0005	M1026

4.4 Configure authselect

Authselect is a utility that simplifies the configuration of user authentication. Authselect offers two ready-made profiles that can be universally used with all modern identity management systems

Authselect makes testing and troubleshooting easy because it only modifies files in these directories:

- `/etc/nsswitch.conf`
- `/etc/pam.d/*` files
- `/etc/dconf/db/distro.d/*` files

You can create and deploy a custom profile by customizing one of the default profiles, the sssd, winbind, or the nis profile. This is particularly useful if Modifying a ready-made authselect profile is not enough for your needs. When you deploy a custom profile, the profile is applied to every user logging into the given host. This would be the recommended method, so that the existing profiles can remain unmodified.

Example of creating a custom authselect profile called custom-profile

```
# authselect create-profile custom-profile -b sssd --symlink-meta
```

WARNING:

Do not use authselect if:

- your host is part of Linux Identity Management. Joining your host to an IdM domain with the `ipa-client-install` command automatically configures SSSD authentication on your host.
- Your host is part of Active Directory via SSSD. Calling the `realm join` command to join your host to an Active Directory domain automatically configures SSSD authentication on your host.
- It is not recommended to change the authselect profiles configured by `ipa-client-install` or `realm join`. If you need to modify them, display the current settings before making any modifications, so you can revert back to them if necessary

4.4.1 Ensure custom authselect profile is used (Manual)

Profile Applicability:

- Level 1 - Server

Description:

A custom profile can be created by copying and customizing one of the default profiles. The default profiles include: sssd, winbind, or the nis. This profile can then be customized to follow site specific requirements.

You can select a profile for the authselect utility for a specific host. The profile will be applied to every user logging into the host.

Rationale:

A custom profile is required to customize many of the pam options.

When you deploy a profile, the profile is applied to every user logging into the given host

Audit:

Ensure authselect is installed

```
# rpm -q authselect
authselect-<version>.amzn2023.0.2.x86_64
```

Run the following command to list the custom profile(s)

```
# authselect list | grep '^-\s*custom'
```

Verify output includes a custom profile:

Example:

```
- custom/custom-profile  Enable SSSD for system authentication (also for
local users only)
```

Run the following command and verify that the current custom authselect profile is in use on the system:

```
# head -1 /etc/authselect/authselect.conf | grep 'custom/'
custom/<CUSTOM_PROFILE_NAME>
```

Remediation:

Run the following command to install authselect if needed

```
# dnf install authselect
```

Run the following command to create a custom authselect profile:

```
# authselect create-profile <custom-profile name> <options>
```

Example:

```
# authselect create-profile custom-profile -b sssd --symlink-meta
```

Run the following command to select a custom authselect profile:

```
# authselect select custom/<CUSTOM PROFILE NAME> {with-<OPTIONS>}
```

Example:

```
# authselect select custom/custom-profile with-sudo with-faillock without-nullok
```

References:

1. authselect(8)
2. NIST SP 800-53 Rev. 5: CA-5

Additional Information:

with the option `--base-on=BASE-ID` or `-b=BASE-ID` the new profile will be based on a profile named BASE-ID.

The base profile location is determined with these steps:

1. If BASE-ID starts with prefix custom/ it is a custom profile.
2. Try if BASE-ID is found in vendor profiles.
3. Try if BASE-ID is found in default profiles.
4. Return an error.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>16.2 Establish and Maintain a Process to Accept and Address Software Vulnerabilities</u> Establish and maintain a process to accept and address reports of software vulnerabilities, including providing a means for external entities to report. The process is to include such items as: a vulnerability handling policy that identifies reporting process, responsible party for handling vulnerability reports, and a process for intake, assignment, remediation, and remediation testing. As part of the process, use a vulnerability tracking system that includes severity ratings, and metrics for measuring timing for identification, analysis, and remediation of vulnerabilities. Review and update documentation annually, or when significant enterprise changes occur that could impact this Safeguard. Third-party application developers need to consider this an externally-facing policy that helps to set expectations for outside stakeholders.			
v7	<u>16.7 Establish Process for Revoking Access</u> Establish and follow an automated process for revoking system access by disabling accounts immediately upon termination or change of responsibilities of an employee or contractor . Disabling these accounts, instead of deleting accounts, allows preservation of audit trails.			

4.4.2 Ensure authselect includes with-faillock (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `pam_faillock.so` module maintains a list of failed authentication attempts per user during a specified interval and locks the account in case there were more than the configured number of consecutive failed authentications (this is defined by the `deny` parameter in the faillock configuration). It stores the failure records into per-user files in the tally directory.

Rationale:

Locking out user IDs after n unsuccessful consecutive login attempts mitigates brute force password attacks against your systems.

Audit:

Run the following commands to verify that faillock is enabled

```
# grep pam_faillock.so /etc/pam.d/password-auth /etc/pam.d/system-auth
```

Output should be similar to:

```
/etc/authselect/password-auth:auth          required
pam_faillock.so preauth silent
/etc/authselect/password-auth:auth          required
pam_faillock.so authfail
/etc/authselect/password-auth:account       required
pam_faillock.so
/etc/authselect/system-auth:auth            required
pam_faillock.so preauth silent
/etc/authselect/system-auth:auth            required
pam_faillock.so authfail
/etc/authselect/system-auth:account         required
pam_faillock.so
```

Remediation:

Run the following commands to include the `with-faillock` option to the current authselect profile:

```
# authselect enable-feature with-faillock
# authselect apply-changes
```

References:

1. faillock(8) - Linux man page
2. pam_faillock(8) - Linux man page
3. NIST SP 800-53 Rev. 5: CM-1,CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.1 Establish and Maintain a Secure Configuration Process</u> Establish and maintain a secure configuration process for enterprise assets (end-user devices, including portable and mobile, non-computing/IoT devices, and servers) and software (operating systems and applications). Review and update documentation annually, or when significant enterprise changes occur that could impact this Safeguard.			
v7	<u>16.7 Establish Process for Revoking Access</u> Establish and follow an automated process for revoking system access by disabling accounts immediately upon termination or change of responsibilities of an employee or contractor . Disabling these accounts, instead of deleting accounts, allows preservation of audit trails.			

4.5 Configure PAM

PAM (Pluggable Authentication Modules) is a service that implements modular authentication modules on UNIX systems. PAM is implemented as a set of shared objects that are loaded and executed when a program needs to authenticate a user. Files for PAM are typically located in the `/etc/pam.d` directory. PAM must be carefully configured to secure system authentication. While this section covers some of PAM, please consult other PAM resources to fully understand the configuration capabilities.

4.5.1 Ensure password creation requirements are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `pam_pwquality.so` module checks the strength of passwords. It performs checks such as making sure a password is not a dictionary word, it is a certain length, contains a mix of characters (e.g. alphabet, numeric, other) and more. The following are definitions of the `pam_pwquality.so` options.

- `try_first_pass` - retrieve the password from a previous stacked PAM module. If not available, then prompt the user for a password.
- `retry=3` - Allow 3 tries before sending back a failure.
- `minlen=14` - password must be 14 characters or more

Either of the following can be used to enforce complex passwords:

- `minclass=4` - provide at least four classes of characters for the new password

OR

- `dcredit=-1` - provide at least one digit
- `ucredit=-1` - provide at least one uppercase character
- `ocredit=-1` - provide at least one special character
- `lcredit=-1` - provide at least one lowercase character

The settings shown above are one possible policy. Alter these values to conform to your own organization's password policies

Rationale:

Strong passwords protect systems from being hacked through brute force methods.

Audit:

Verify password creation requirements conform to organization policy:

Run the following command and verify that retry conforms to organization policy.

```
# grep pam_pwquality.so /etc/pam.d/system-auth /etc/pam.d/password-auth
```

Output should be similar to:

```
/etc/pam.d/system-auth:password requisite pam_pwquality.so try_first_pass  
local_users_only enforce_for_root retry=3  
/etc/pam.d/password-auth:password requisite pam_pwquality.so try_first_pass  
local_users_only enforce_for_root retry=3
```

Run the following commands and verify password length requirements conform to organization policy.

```
# grep ^minlen /etc/security/pwquality.conf
```

Verify minlen is 14 or more

Run one of the following commands and verify that password complexity conforms to organization policy.

```
# grep ^minclass /etc/security/pwquality.conf
```

OR

```
# grep -E "^\\s*\\Scredit\\s*" /etc/security/pwquality.conf
```

Remediation:

Edit the file `/etc/security/pwquality.conf` and add or modify the following line for password length to conform to site policy

```
minlen = 14
```

Edit the file `/etc/security/pwquality.conf` and add or modify the following line for password complexity to conform to site policy

```
minclass = 4
```

OR

```
dcredit = -1
ucredit = -1
ocredit = -1
lcredit = -1
```

Run the following script to update the system-auth and password-auth files

```
{
    #!/usr/bin/env bash

    for fn in system-auth password-auth; do
        file="/etc/authselect/${head -1 /etc/authselect/authselect.conf | grep
'custom/'}/$fn"
        if ! grep -Pq --
'^\h*password\h+requisite\h+pam_pwquality.so(\h+[^#\n\r]+)?\h+.*enforce_for_r
oot\b.*$' "$file"; then
            sed -ri 's/^\s*(password\s+requisite\s+pam_pwquality.so\s+)(.*)$/\1\2
enforce_for_root/' "$file"
        fi
        if grep -Pq --
'^\h*password\h+requisite\h+pam_pwquality.so(\h+[^#\n\r]+)?\h+retry=([4-
9]|[1-9][0-9]+)\b.*$' "$file"; then
            sed -ri '/pwquality/s/retry=\S+/retry=3/' "$file"
        elif ! grep -Pq --
'^\h*password\h+requisite\h+pam_pwquality.so(\h+[^#\n\r]+)?\h+retry=\d+\b.*$'
"$file"; then
            sed -ri 's/^\s*(password\s+requisite\s+pam_pwquality.so\s+)(.*)$/\1\2
retry=3/' "$file"
        fi
    done
    authselect apply-changes
}
```

References:

1. NIST SP 800-53 Rev. 5: IA-5

Additional Information:

all default authselect profiles have `pam_pwquality` enabled with the expectation that options will be specified in `pwquality.conf`

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.2 Use Unique Passwords Use unique passwords for all enterprise assets. Best practice implementation includes, at a minimum, an 8-character password for accounts using MFA and a 14-character password for accounts not using MFA.			
v7	4.4 Use Unique Passwords Where multi-factor authentication is not supported (such as local administrator, root, or service accounts), accounts will use passwords that are unique to that system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1110, T1110.001, T1110.002, T1110.003, T1178.001, T1178.002, T1178.003, T1178.004	TA0006	M1027

4.5.2 Ensure lockout for failed password attempts is configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Lock out users after *n* unsuccessful consecutive login attempts.

- `deny=<n>` - Number of attempts before the account is locked
- `unlock_time=<n>` - Time in seconds before the account is unlocked

Note: The maximum configurable value for `unlock_time` is 604800

Rationale:

Locking out user IDs after *n* unsuccessful consecutive login attempts mitigates brute force password attacks against your systems.

Impact:

Use of `unlock_time=0` may allow an attacker to cause denial of service to legitimate users.

Audit:

Verify password lockouts are configured. Depending on the version you are running, follow **one** of the two methods bellow.

- `deny` **should not** be 0 (never) or greater than 5
- `unlock_time` **should** be 0 (never) or 900 seconds or more.

Run the following command to verify that Number of failed logon attempts before the account is locked is no greater than 5:

```
# grep -E '^\\s*deny\\s*=\\s*[1-5]\\b' /etc/security/faillock.conf
deny = 5
```

Run the following command to verify that the time in seconds before the account is unlocked is either 0 (never) or 900 or more.

```
# grep -E '^\\s*unlock_time\\s*=\\s*(0|9[0-9][0-9]|[1-9][0-9][0-9][0-9]+)\\b' /etc/security/faillock.conf
unlock_time = 900
```

Remediation:

Set password lockouts and unlock times to conform to site policy. `deny` should be greater than 0 and no greater than 5. `unlock_time` should be 0 (never), or 900 seconds or greater.

Edit `/etc/security/faillock.conf` and update or add the following lines:

```
deny = 5
unlock_time = 900
```

Default Value:

`deny = 3`

`unlock_time = 600`

References:

1. NIST SP 800-53 Rev. 5: AC-1, AC-2

Additional Information:

Additional module options may be set, recommendation only covers those listed here.

If a user has been locked out because they have reached the maximum consecutive failure count defined by `deny=` in the `pam_faillock.so` module, the user can be unlocked by issuing the command `faillock --user <USERNAME> --reset`. This command sets the failed count to 0, effectively unlocking the user.

Use of the "audit" keyword may log credentials in the case of user error during authentication. This risk should be evaluated in the context of the site policies of your organization.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	6.2 Establish an Access Revoking Process Establish and follow a process, preferably automated, for revoking access to enterprise assets, through disabling accounts immediately upon termination, rights revocation, or role change of a user. Disabling accounts, instead of deleting accounts, may be necessary to preserve audit trails.			
v7	16.7 Establish Process for Revoking Access Establish and follow an automated process for revoking system access by disabling accounts immediately upon termination or change of responsibilities of an employee or contractor. Disabling these accounts, instead of deleting accounts, allows preservation of audit trails.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1110, T1110.001, T1110.003	TA0006	M1027

4.5.3 Ensure password reuse is limited (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `/etc/security/opasswd` file stores the users' old passwords and can be checked to ensure that users are not recycling recent passwords.

- `remember=<5>` - Number of old passwords to remember

Rationale:

Forcing users not to reuse their past 5 passwords make it less likely that an attacker will be able to guess the password.

Note: These change only apply to accounts configured on the local system.

Audit:

Run the following command and verify that the remembered password history is 5 or more

```
# grep -P
'^\h*password\h+(requisite|sufficient)\h+(pam_pwhistory\.so|pam_unix\.so)\h+([^\n\r]+\h+)?remember=([5-9]|[1-9][0-9]+)\h*(\h+.*?)?$' /etc/pam.d/system-auth
```

The output should be similar to:

```
password      requisite      pam_pwhistory.so try_first_pass local_users_only
enforce_for_root retry=3 remember=5
password      sufficient    pam_unix.so sha512 shadow try_first_pass
use_authtok remember=5
```

Remediation:

Set remembered password history to conform to site policy.

Run the following script to add or modify the `pam_pwhistory.so` and `pam_unix.so` lines to include the `remember` option:

```
#!/usr/bin/env bash

{
  file="/etc/authselect/$(head -1 /etc/authselect/authselect.conf | grep
'custom/')/system-auth"
  if ! grep -Pq --
'^\h*password\h+(requisite|required|sufficient)\h+pam_pwhistory\.so\h+([\^#\n\r]+\h+)?remember=([5-9]|[1-9][0-9]+)\b.*$' "$file"; then
    if grep -Pq --
'^\h*password\h+(requisite|required|sufficient)\h+pam_pwhistory\.so\h+([\^#\n\r]+\h+)?remember=\d+\b.*$' "$file"; then
      sed -ri
's/^\s*(password\s+(requisite|required|sufficient)\s+pam_pwhistory\.so\s+([\^#\n\r]+\s+)?)(remember=\S+\s*)(\s+.*?)?$/\1 remember=5 \5/' $file
    elif grep -Pq --
'^\h*password\h+(requisite|required|sufficient)\h+pam_pwhistory\.so\h+([\^#\n\r]+\h+)?.*$' "$file"; then
      sed -ri
'/^\s*password\s+(requisite|required|sufficient)\s+pam_pwhistory\.so/ s/$/
remember=5/' $file
    else
      sed -ri
'/^\s*password\s+(requisite|required|sufficient)\s+pam_unix\.so/i password
required      pam_pwhistory.so  remember=5 use_authok' $file
    fi
  fi
  if ! grep -Pq --
'^\h*password\h+(requisite|required|sufficient)\h+pam_unix\.so\h+([\^#\n\r]+\h+)?remember=([5-9]|[1-9][0-9]+)\b.*$' "$file"; then
    if grep -Pq --
'^\h*password\h+(requisite|required|sufficient)\h+pam_unix\.so\h+([\^#\n\r]+\h+)?remember=\d+\b.*$' "$file"; then
      sed -ri
's/^\s*(password\s+(requisite|required|sufficient)\s+pam_unix\.so\s+([\^#\n\r]+\s+)?)(remember=\S+\s*)(\s+.*?)?$/\1 remember=5 \5/' $file
    else
      sed -ri
'/^\s*password\s+(requisite|required|sufficient)\s+pam_unix\.so/ s/$/
remember=5/' $file
    fi
  fi
  authselect apply-changes
}
```

References:

1. NIST SP 800-53 Rev. 5: IA-5

Additional Information:

Additional module options may be set, recommendation only covers those listed here.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.2 Use Unique Passwords Use unique passwords for all enterprise assets. Best practice implementation includes, at a minimum, an 8-character password for accounts using MFA and a 14-character password for accounts not using MFA.			
v7	4.4 Use Unique Passwords Where multi-factor authentication is not supported (such as local administrator, root, or service accounts), accounts will use passwords that are unique to that system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1078, T1078.001, T1078.002, T1078.003, T1078.004, T1110, T1110.004	TA0006	M1027

4.5.4 Ensure password hashing algorithm is SHA-512 (Automated)

Profile Applicability:

- Level 1 - Server

Description:

A cryptographic hash function converts an arbitrary-length input into a fixed length output. Password hashing performs a one-way transformation of a password, turning the password into another string, called the hashed password.

Rationale:

The SHA-512 algorithm provides stronger hashing than other hashing algorithms used for password hashing with Linux, providing additional protection to the system by increasing the level of effort for an attacker to successfully determine passwords.

Note: These changes only apply to accounts configured on the local system.

Audit:

Verify password hashing algorithm is sha512 or yescrypt:

Run the following command to verify the hashing algorithm is sha512 or yescrypt in /etc/libuser.conf:

```
# grep -Ei '^s*crypt_style\s*=\s*(sha512|yescrypt)\b' /etc/libuser.conf
crypt_style = sha512
```

Run the following command to verify the hashing algorithm is sha512 or yescrypt in /etc/login.defs:

```
# grep -Ei '^s*ENCRYPT_METHOD\s+(SHA512|yescrypt)\b' /etc/login.defs
ENCRYPT_METHOD SHA512

-OR-

ENCRYPT_METHOD YESCRYPT
```

Run the following command to verify the hashing algorithm is configured with pam_unix.so in /etc/pam.d/system-auth and /etc/pam.d/password-auth:

```
# grep -P --
'^\h*password\h+(requisite|required)sufficient)\h+pam_unix\.so(\h+[\h*\n\r]+)?
\h+(sha512|yescrypt)\b.*$' /etc/pam.d/password-auth /etc/pam.d/system-auth
```

The output should be similar to:

```
/etc/pam.d/password-auth:password    sufficient    pam_unix.so sha512 shadow
try_first_pass use_authtok remember=5
/etc/pam.d/system-auth:password      sufficient    pam_unix.so sha512 shadow
try_first_pass use_authtok remember=5
```

Remediation:

Set password hashing algorithm to sha512.

Edit `/etc/libuser.conf` and edit or add the following line:

```
crypt_style = sha512
```

Edit `/etc/login.defs` and edit or add the following line:

```
ENCRYPT_METHOD SHA512
```

```
-OR-
```

```
ENCRYPT_METHOD YESCRYPT
```

Run the following script to configure `pam_unix.so` to use the sha512 hashing algorithm:

```
#!/usr/bin/env bash

{
    for fn in system-auth password-auth; do
        file="/etc/authselect/${head -1 /etc/authselect/authselect.conf | grep
'custom/'}/$fn"
        if ! grep -Pq --
'^\h*password\h+(requisite|required|sufficient)\h+pam_unix\.so(\h+[^#\n\r]+)?
\h+sha512\b.*$' "$file"; then
            if grep -Pq --
'^\h*password\h+(requisite|required|sufficient)\h+pam_unix\.so(\h+[^#\n\r]+)?
\h+(md5|blowfish|bigcrypt|sha256|yescrypt)\b.*$' "$file"; then
                sed -ri 's/(md5|blowfish|bigcrypt|sha256|yescrypt)/sha512/' "$file"
            else
                sed -ri
's/(\^s*password\s+(requisite|required|sufficient)\s+pam_unix.so\s+)(.*)$/\1s
ha512 \3/' "$file"
            fi
        fi
    done
    authselect apply-changes
}
```

Note: This only effects local users and passwords created after updating the files to use sha512. If it is determined that the password algorithm being used is not SHA-512, once it is changed, it is recommended that all user ID's be immediately expired and forced to change their passwords on next login.

References:

1. NIST SP 800-53 Rev. 5: IA-5

Additional Information:

Additional module options may be set, recommendation only covers those listed here.

The following command may be used to expire all non-system user ID's immediately and force them to change their passwords on next login. Any system accounts that need to be expired should be carefully done separately by the system administrator to prevent any potential problems.

```
# awk -F: ' ( $3<"$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)"' && $1 != "nfsnobody" ) { print $1 }' /etc/passwd | xargs -n 1 chage -d 0
```

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.11 <u>Encrypt Sensitive Data at Rest</u> Encrypt sensitive data at rest on servers, applications, and databases containing sensitive data. Storage-layer encryption, also known as server-side encryption, meets the minimum requirement of this Safeguard. Additional encryption methods may include application-layer encryption, also known as client-side encryption, where access to the data storage device(s) does not permit access to the plain-text data.			
v7	16.4 <u>Encrypt or Hash all Authentication Credentials</u> Encrypt or hash with a salt all authentication credentials when stored.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1003, T1003.008, T1110, T1110.002	TA0006	M1041

4.6 User Accounts and Environment

This section provides guidance on setting up secure defaults for system and user accounts and their environment.

4.6.1 Set Shadow Password Suite Parameters

While a majority of the password control parameters have been moved to PAM, some parameters are still available through the shadow password suite. Any changes made to `/etc/login.defs` will only be applied if the `usermod` command is used. If user IDs are added a different way, use the `chage` command to effect changes to individual user IDs.

4.6.1.1 Ensure password expiration is 365 days or less (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `PASS_MAX_DAYS` parameter in `/etc/login.defs` allows an administrator to force passwords to expire once they reach a defined age. It is recommended that the `PASS_MAX_DAYS` parameter be set to less than or equal to 365 days.

Rationale:

The window of opportunity for an attacker to leverage compromised credentials or successfully compromise credentials via an online brute force attack is limited by the age of the password. Therefore, reducing the maximum age of a password also reduces an attacker's window of opportunity.

Audit:

Run the following command and verify `PASS_MAX_DAYS` conforms to site policy (no more than 365 days):

```
# grep PASS_MAX_DAYS /etc/login.defs
PASS_MAX_DAYS 365
```

Run the following command and Review list of users and `PASS_MAX_DAYS` to verify that all users' `PASS_MAX_DAYS` conforms to site policy (no more than 365 days):

```
# grep -E '^[^:]+:[^!]*' /etc/shadow | cut -d: -f1,5
<user>:<PASS_MAX_DAYS>
```

Remediation:

Set the `PASS_MAX_DAYS` parameter to conform to site policy in `/etc/login.defs` :

```
PASS_MAX_DAYS 365
```

Modify user parameters for all users with a password set to match:

```
# chage --maxdays 365 <user>
```

References:

1. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

Additional Information:

You can also check this setting in `/etc/shadow` directly. The 5th field should be 365 or less for all users with a password.

Note: A value of -1 will disable password expiration. Additionally, the password expiration must be greater than the minimum days between password changes or users will be unable to change their password.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.1 <u>Establish and Maintain a Secure Configuration Process</u> Establish and maintain a secure configuration process for enterprise assets (end-user devices, including portable and mobile, non-computing/IoT devices, and servers) and software (operating systems and applications). Review and update documentation annually, or when significant enterprise changes occur that could impact this Safeguard.			
v7	4.4 <u>Use Unique Passwords</u> Where multi-factor authentication is not supported (such as local administrator, root, or service accounts), accounts will use passwords that are unique to that system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1078, T1078.001, T1078.002, T1078.003, T1078.004, T1110, T1110.001, T1110.002, T1110.003, T1110.004	TA0001	M1027

4.6.1.2 Ensure minimum days between password changes is configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `PASS_MIN_DAYS` parameter in `/etc/login.defs` allows an administrator to prevent users from changing their password until a minimum number of days have passed since the last time the user changed their password. It is recommended that `PASS_MIN_DAYS` parameter be set to 1 or more days.

Rationale:

By restricting the frequency of password changes, an administrator can prevent users from repeatedly changing their password in an attempt to circumvent password reuse controls.

Audit:

Run the following command and verify `PASS_MIN_DAYS` conforms to site policy (no less than 1 day):

```
# grep PASS_MIN_DAYS /etc/login.defs  
  
PASS_MIN_DAYS 1
```

Run the following command and Review list of users and `PAS_MIN_DAYS` to Verify that all users' `PAS_MIN_DAYS` conforms to site policy (no less than 1 day):

```
# awk -F : '(!/^[:]+:[^!*/] / && $4 < 1){print $1 " " $4}' /etc/shadow  
  
No <user>:<PASS_MIN_DAYS> should be returned
```

Remediation:

Set the `PASS_MIN_DAYS` parameter to 1 in `/etc/login.defs`:

```
PASS_MIN_DAYS 1
```

Modify user parameters for all users with a password set to match:

```
# chage --mindays 1 <user>
```

Default Value:

`PASS_MIN_DAYS 0`

References:

1. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

Additional Information:

You can also check this setting in `/etc/shadow` directly. The 4th field should be 1 or more for all users with a password.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.2 Use Unique Passwords Use unique passwords for all enterprise assets. Best practice implementation includes, at a minimum, an 8-character password for accounts using MFA and a 14-character password for accounts not using MFA.			
v7	4.4 Use Unique Passwords Where multi-factor authentication is not supported (such as local administrator, root, or service accounts), accounts will use passwords that are unique to that system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1078, T1078.001, T1078.002, T1078.003, T1078.004, T1110.004	TA0006	M1027

4.6.1.3 Ensure password expiration warning days is 7 or more (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `PASS_WARN_AGE` parameter in `/etc/login.defs` allows an administrator to notify users that their password will expire in a defined number of days. It is recommended that the `PASS_WARN_AGE` parameter be set to 7 or more days.

Rationale:

Providing an advance warning that a password will be expiring gives users time to think of a secure password. Users caught unaware may choose a simple password or write it down where it may be discovered.

Audit:

Run the following command and verify `PASS_WARN_AGE` conforms to site policy (No less than 7 days):

```
# grep PASS_WARN_AGE /etc/login.defs
PASS_WARN_AGE 7
```

Verify all users with a password have their number of days of warning before password expires set to 7 or more:

Run the following command and Review list of users and `PASS_WARN_AGE` to verify that all users' `PASS_WARN_AGE` conforms to site policy (No less than 7 days):

```
# grep -E ^[^:]+:[^\!*] /etc/shadow | cut -d: -f1,6
<user>:<PASS_WARN_AGE>
```

Remediation:

Set the `PASS_WARN_AGE` parameter to 7 in `/etc/login.defs`:

```
PASS_WARN_AGE 7
```

Modify user parameters for all users with a password set to match:

```
# chage --warndays 7 <user>
```

References:

1. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

Additional Information:

You can also check this setting in `/etc/shadow` directly. The 6th field should be 7 or more for all users with a password.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.1 Establish and Maintain a Secure Configuration Process Establish and maintain a secure configuration process for enterprise assets (end-user devices, including portable and mobile, non-computing/IoT devices, and servers) and software (operating systems and applications). Review and update documentation annually, or when significant enterprise changes occur that could impact this Safeguard.			
v7	4.4 Use Unique Passwords Where multi-factor authentication is not supported (such as local administrator, root, or service accounts), accounts will use passwords that are unique to that system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1078	TA0006	M1027

References:

1. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

Additional Information:

You can also check this setting in `/etc/shadow` directly. The 7th field should be 30 or less for all users with a password.

Note: A value of -1 would disable this setting.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.2 Use Unique Passwords Use unique passwords for all enterprise assets. Best practice implementation includes, at a minimum, an 8-character password for accounts using MFA and a 14-character password for accounts not using MFA.			
v7	4.4 Use Unique Passwords Where multi-factor authentication is not supported (such as local administrator, root, or service accounts), accounts will use passwords that are unique to that system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1078, T1078.002, T1078.003	TA0001	M1027

4.6.1.5 Ensure all users last password change date is in the past (Automated)

Profile Applicability:

- Level 1 - Server

Description:

All users should have a password change date in the past.

Rationale:

If a user's recorded password change date is in the future, then they could bypass any set password expiration.

Audit:

Run the following command and verify nothing is returned

```
# awk -F: '/^[^:]+:[^!*/]{print $1}' /etc/shadow | while read -r usr; \
do change=$(date -d "$(chage --list $usr | grep '^Last password change' | cut
-d: -f2 | grep -v 'never$')" +%s); \
if [[ "$change" -gt "$(date +%s)" ]]; then \
echo "User: \"$usr\" last password change was \"$$(chage --list $usr | grep
'^Last password change' | cut -d: -f2)\""; fi; done
```

Remediation:

Investigate any users with a password change date in the future and correct them. Locking the account, expiring the password, or resetting the password manually may be appropriate.

References:

1. NIST SP 800-53 Rev. 5: IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.2 Use Unique Passwords Use unique passwords for all enterprise assets. Best practice implementation includes, at a minimum, an 8-character password for accounts using MFA and a 14-character password for accounts not using MFA.			
v7	4.4 Use Unique Passwords Where multi-factor authentication is not supported (such as local administrator, root, or service accounts), accounts will use passwords that are unique to that system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1078, T1078.001, T1078.002, T1078.003, T1078.004, T1110, T1110.001, T1110.002, T1110.003, T1110.004	TA0006	M1027

4.6.2 Ensure system accounts are secured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

There are a number of accounts provided with most distributions that are used to manage applications and are not intended to provide an interactive shell. Furthermore, a user may add special accounts that are not intended to provide an interactive shell.

Rationale:

It is important to make sure that accounts that are not being used by regular users are prevented from being used to provide an interactive shell. By default, most distributions set the password field for these accounts to an invalid string, but it is also recommended that the shell field in the password file be set to the `nologin` shell. This prevents the account from potentially being used to run any commands.

Audit:

System accounts

Check critical system accounts for `nologin`

Run the following command:

```
# awk -F: '($1!~/^(root|halt|sync|shutdown|nfsnobody)$/ && ($3<"$(awk '/^\\s*UID_MIN/{print $2}' /etc/login.defs)" || $3 == 65534) && $7!~/^(\\usr)?\\sbin\\nologin$/ ) { print $1 }' /etc/passwd
```

Verify no results are returned.

Disabled accounts

Ensure all accounts that configured the shell as `nologin` also have their passwords disabled.

Run the following command:

```
# awk -F: '/nologin/ {print $1}' /etc/passwd | xargs -I '{}' passwd -S '{}' | awk '($2!="L" && $2!="LK") {print $1}'
```

Verify no results are returned.

Remediation:

System accounts

Set the shell for any accounts returned by the audit to `nologin`:

```
# usermod -s $(command -v nologin) <user>
```

Disabled accounts

Lock any non root accounts returned by the audit:

```
# usermod -L <user>
```

Large scale changes

The following command will set all system accounts to `nologin`:

```
# awk -F: '($1!~/^(root|halt|sync|shutdown|nfsnobody)$/ && ($3<"$(awk '/^s*UID_MIN/{print $2}' /etc/login.defs)" || $3 == 65534)) { print $1 }' /etc/passwd | while read user; do usermod -s $(command -v nologin) $user >/dev/null; done
```

The following command will automatically lock all accounts that have their shell set to `nologin`:

```
# awk -F: '/nologin/ {print $1}' /etc/passwd | while read user; do usermod -L $user; done
```

References:

1. NIST SP 800-53 Rev. 5: AC-2, AC-3, AC-5, MP-2

Additional Information:

The `root`, `sync`, `shutdown`, and `halt` users are exempted from requiring a non-login shell.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1078, T1078.001, T1078.003	TA0005	M1026

4.6.3 Ensure default user shell timeout is 900 seconds or less (Automated)

Profile Applicability:

- Level 1 - Server

Description:

`TMOUT` is an environmental setting that determines the timeout of a shell in seconds.

- `TMOUT=n` - Sets the shell timeout to *n* seconds. A setting of `TMOUT=0` disables timeout.
- `readonly TMOUT` - Sets the `TMOUT` environmental variable as readonly, preventing unwanted modification during run-time.
- `export TMOUT` - exports the `TMOUT` variable

System Wide Shell Configuration Files:

- `/etc/profile` - used to set system wide environmental variables on users shells. The variables are sometimes the same ones that are in the `.bash_profile`, however this file is used to set an initial `PATH` or `PS1` for all shell users of the system. **is only executed for interactive *login* shells, or shells executed with the `--login` parameter.**
- `/etc/profile.d` - `/etc/profile` will execute the scripts within `/etc/profile.d/*.sh`. It is recommended to place your configuration in a shell script within `/etc/profile.d` to set your own system wide environmental variables.
- `/etc/bashrc` - System wide version of `.bashrc`. In Fedora derived distributions, `/etc/bashrc` also invokes `/etc/profile.d/*.sh` if *non-login* shell, but redirects output to `/dev/null` if *non-interactive*. **Is only executed for *interactive* shells or if `BASH_ENV` is set to `/etc/bashrc`.**

Rationale:

Setting a timeout value reduces the window of opportunity for unauthorized user access to another user's shell session that has been left unattended. It also ends the inactive session and releases the resources associated with that session.

Audit:

Run the following script to verify that `TMOUT` is configured to: include a timeout of no more than 900 seconds, to be `readonly`, to be `exported`, and is not being changed to a longer timeout.

```
#!/usr/bin/env bash

{
    output1="" output2=""
    [ -f /etc/bashrc ] && BRC="/etc/bashrc"
    for f in "$BRC" /etc/profile /etc/profile.d/*.sh ; do
        grep -Pq '^s*([^\#]+\s+)?TMOUT=(900|[1-8][0-9][0-9]|[1-9][0-9]|[1-9])\b' "$f" && grep -Pq '^s*([^\#]+\s+)?readonly\s+TMOUT(\s+|\s*;\s*\$|=900|[1-8][0-9][0-9]|[1-9][0-9]|[1-9]))\b' "$f" && grep -Pq '^s*([^\#]+\s+)?export\s+TMOUT(\s+|\s*;\s*\$|=900|[1-8][0-9][0-9]|[1-9][0-9]|[1-9]))\b' "$f" &&
        output1="$f"
    done
    grep -Pq '^s*([^\#]+\s+)?TMOUT=(9[0-9][1-9]|9[1-9][0-9]|0+|[1-9]\d{3,})\b' /etc/profile /etc/profile.d/*.sh "$BRC" && output2=$(grep -Ps '^s*([^\#]+\s+)?TMOUT=(9[0-9][1-9]|9[1-9][0-9]|0+|[1-9]\d{3,})\b' /etc/profile /etc/profile.d/*.sh $BRC)
    if [ -n "$output1" ] && [ -z "$output2" ]; then
        echo -e "\nPASSED\n\nTMOUT is configured in: \"$output1\"\n"
    else
        [ -z "$output1" ] && echo -e "\nFAILED\n\nTMOUT is not configured\n"
        [ -n "$output2" ] && echo -e "\nFAILED\n\nTMOUT is incorrectly configured in: \"$output2\"\n"
    fi
}
```


Remediation:

Review `/etc/bashrc`, `/etc/profile`, and all files ending in `*.sh` in the `/etc/profile.d/` directory and remove or edit all `TMOUT=_n_` entries to follow local site policy. `TMOUT` should not exceed 900 or be equal to 0.

Configure `TMOUT` in **one** of the following files:

- A file in the `/etc/profile.d/` directory ending in `.sh`
- `/etc/profile`
- `/etc/bashrc`

TMOUT configuration examples:

- As multiple lines:

```
TMOUT=900
readonly TMOUT
export TMOUT
```

- As a single line:

```
readonly TMOUT=900 ; export TMOUT
```

References:

1. NIST SP 800-53 Rev. 5: AC-11

Additional Information:

The audit and remediation in this recommendation apply to bash and shell. If other shells are supported on the system, it is recommended that their configuration files also are checked. Other methods of setting a timeout exist for other shells not covered here.

Ensure that the timeout conforms to your local policy.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.3 <u>Configure Automatic Session Locking on Enterprise Assets</u> Configure automatic session locking on enterprise assets after a defined period of inactivity. For general purpose operating systems, the period must not exceed 15 minutes. For mobile end-user devices, the period must not exceed 2 minutes.			
v7	16.11 <u>Lock Workstation Sessions After Inactivity</u> Automatically lock workstation sessions after a standard period of inactivity.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1078	TA0005	M1026

4.6.4 Ensure default group for the root account is GID 0 (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `usermod` command can be used to specify which group the `root` account belongs to. This affects permissions of files that are created by the `root` account.

Rationale:

Using GID 0 for the `root` account helps prevent `root` -owned files from accidentally becoming accessible to non-privileged users.

Audit:

Run the following command and verify the result is 0 :

```
# grep "^root:" /etc/passwd | cut -f4 -d:
0
```

Remediation:

Run the following command to set the `root` account default group to GID 0 :

```
# usermod -g 0 root
```

References:

1. NIST SP 800-53 Rev. 5: CM-1,CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.1 Establish and Maintain a Secure Configuration Process Establish and maintain a secure configuration process for enterprise assets (end-user devices, including portable and mobile, non-computing/IoT devices, and servers) and software (operating systems and applications). Review and update documentation annually, or when significant enterprise changes occur that could impact this Safeguard.			
v7	5.1 Establish Secure Configurations Maintain documented, standard security configuration standards for all authorized operating systems and software.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1548, T1548.000	TA0005	M1026

4.6.5 Ensure default user umask is 027 or more restrictive (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The user file-creation mode mask (`umask`) is used to determine the file permission for newly created directories and files. In Linux, the default permissions for any newly created directory is 0777 (`rw-rw-rw-`), and for any newly created file it is 0666 (`rw-rw-rw-`). The `umask` modifies the default Linux permissions by restricting (masking) these permissions. The `umask` is not simply subtracted, but is processed bitwise. Bits set in the `umask` are cleared in the resulting file mode.

`umask` can be set with either `octal` or `Symbolic` values:

- **Octal (Numeric) Value** - Represented by either three or four digits. ie `umask 0027` or `umask 027`. If a four digit `umask` is used, the first digit is ignored. The remaining three digits effect the resulting permissions for user, group, and world/other respectively.
- **Symbolic Value** - Represented by a comma separated list for User `u`, group `g`, and world/other `o`. The permissions listed are not masked by `umask`. ie a `umask` set by `umask u=rwx,g=rx,o=` is the `Symbolic` equivalent of the `Octal umask 027`. This `umask` would set a newly created directory with file mode `drwxr-x---` and a newly created file with file mode `rw-r-----`.

The default `umask` can be set to use the `pam_umask` module or in a `System Wide Shell Configuration File`. The user creating the directories or files has the discretion of changing the permissions via the `chmod` command, or choosing a different default `umask` by adding the `umask` command into a `User Shell Configuration File`, (`.bash_profile` or `.bashrc`), in their home directory.

Setting the default umask:

- `pam_umask` module:
 - will set the `umask` according to the system default in `/etc/login.defs` and user settings, solving the problem of different `umask` settings with different shells, display managers, remote sessions etc.
 - `umask=<mask>` value in the `/etc/login.defs` file is interpreted as Octal
 - Setting `USERGROUPS_ENAB` to `yes` in `/etc/login.defs` (default):
 - will enable setting of the `umask` group bits to be the same as owner bits. (examples: `022` -> `002`, `077` -> `007`) for non-root users, if the `uid` is the same as `gid`, and `username` is the same as the `<primary group name>`
 - `userdel` will remove the user's group if it contains no more members, and `useradd` will create by default a group with the name of the user
- System Wide Shell Configuration File:
 - `/etc/profile` - used to set system wide environmental variables on users shells. The variables are sometimes the same ones that are in the `.bash_profile`, however this file is used to set an initial `PATH` or `PS1` for all shell users of the system. **is only executed for interactive login shells, or shells executed with the `--login` parameter.**
 - `/etc/profile.d` - `/etc/profile` will execute the scripts within `/etc/profile.d/*.sh`. It is recommended to place your configuration in a shell script within `/etc/profile.d` to set your own system wide environmental variables.
 - `/etc/bashrc` - System wide version of `.bashrc`. In Fedora derived distributions, `etc/bashrc` also invokes `/etc/profile.d/*.sh` if *non-login* shell, but redirects output to `/dev/null` if *non-interactive*. **Is only executed for interactive shells or if `BASH_ENV` is set to `/etc/bashrc`.**

User Shell Configuration Files:

- `~/.bash_profile` - Is executed to configure your shell before the initial command prompt. **Is only read by login shells.**
- `~/.bashrc` - Is executed for interactive shells. **only read by a shell that's both interactive and non-login**

Rationale:

Setting a secure default value for `umask` ensures that users make a conscious choice about their file permissions. A permissive `umask` value could result in directories or files with excessive permissions that can be read and/or written to by unauthorized users.

Audit:

Run the following to verify:

- A default user `umask` is set to enforce a newly created directories' permissions to be 750 (`drwxr-x---`), and a newly created file's permissions be 640 (`rw-r-----`), or more restrictive
- No less restrictive System Wide `umask` is set

Run the following script to verify that a default user `umask` is set enforcing a newly created directories's permissions to be 750 (`drwxr-x---`), and a newly created file's permissions be 640 (`rw-r-----`), or more restrictive:

```
#!/bin/bash

{
    passing=""
    grep -Eiq '^\s*UMASK\s+([0-7][2-7]7|[0-7][2-7]7)\b' /etc/login.defs &&
    grep -Eiq '^\s*USERGROUPS_ENAB\s*"no"?'\b' /etc/login.defs && grep -Eq
    '^\s*session\s+(optional|required)\s+pam_umask\.so\b'
    /etc/pam.d/common-session && passing=true
    grep -REiq '^\s*UMASK\s+\s*([0-7][2-7]7|[0-7][2-
7]7|u=(r?w?|x?) (r?w?|x?) (r?w?|x?),g=(r?x?|x?r?),o=)\b' /etc/profile*
    /etc/bashrc* && passing=true
    [ "$passing" = true ] && echo "Default user umask is set"
}
```

Verify output is: "Default user umask is set"

Run the following to verify that no less restrictive system wide `umask` is set:

```
# grep -RPi '^(^|^[^#]*)\s*umask\s+([0-7][0-7][01][0-7]\b|[0-7][0-7][0-7][0-
6]\b|[0-7][01][0-7]\b|[0-7][0-7][0-
6]\b|(u=[rx]{0,3},)?(g=[rx]{0,3},)?o=[rx]+\b|(u=[rx]{1,3},)?g=[^rx]{1,3}(
,o=[rx]{0,3})?\b)' /etc/login.defs /etc/profile* /etc/bashrc*

No file should be returned
```

Remediation:

Review `/etc/bashrc`, `/etc/profile`, and all files ending in `.sh` in the `/etc/profile.d/` directory and remove or edit all `umask` entries to follow local site policy. Any remaining entries should be: `umask 027`, `umask u=rwx,g=rx,o=` or more restrictive.

Configure `umask` in **one** of the following files:

- A file in the `/etc/profile.d/` directory ending in `.sh`
- `/etc/profile`
- `/etc/bashrc`

Example:

```
# vi /etc/profile.d/set_umask.sh

umask 027
```

Run the following command and remove or modify the `umask` of any returned files:

```
# grep -RPi ' (^|^ [^#]*) \s*umask\s+ ([0-7] [0-7] [01] [0-7] \b| [0-7] [0-7] [0-7] [0-7] \b| [0-7] [01] [0-7] \b| [0-7] [0-7] [0-7] [0-7] \b| (u=[rxw]{0,3},)? (g=[rxw]{0,3},)? o=[rxw]+\b| (u=[rxw]{1,3},)? g=[^rx]{1,3} (,o=[rxw]{0,3})?\b)' /etc/login.defs /etc/profile* /etc/bashrc*
```

Follow one of the following methods to set the default user `umask`:

Edit `/etc/login.defs` and edit the `UMASK` and `USERGROUPS_ENAB` lines as follows:

```
UMASK 027

USERGROUPS_ENAB no
```

Edit the files `/etc/pam.d/password-auth` and `/etc/pam.d/system-auth` and add or edit the following:

```
session      optional      pam_umask.so
```

OR Configure `umask` in one of the following files:

- A file in the `/etc/profile.d/` directory ending in `.sh`
- `/etc/profile`
- `/etc/bashrc`

Example: `/etc/profile.d/set_umask.sh`

```
umask 027
```

Note: this method only applies to `bash` and `shell`. If other shells are supported on the system, it is recommended that their configuration files also are checked.

Default Value:

UMASK 022

References:

1. NIST SP 800-53 Rev. 5: AC-3, MP-2

Additional Information:

- Other methods of setting a default user umask exist
- If other methods are in use in your environment they should be audited
- The default user umask can be overridden with a user specific umask
- The user creating the directories or files has the discretion of changing the permissions:
 - Using the chmod command
 - Setting a different default umask by adding the umask command into a User Shell Configuration File, (.bashrc), in their home directory
 - Manually changing the umask for the duration of a login session by running the umask command

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1083	TA0007	

4.6.6 Ensure root password is set (Manual)

Profile Applicability:

- Level 1 - Server

Description:

There are a number of methods to access the root account directly. Without a password set any user would be able to gain access and thus control over the entire system.

Rationale:

Access to `root` should be secured at all times.

Impact:

If there are any automated processes that relies on access to the root account without authentication, they will fail after remediation.

Audit:

Run the following command:

```
# passwd -S root
```

Verify that the output contains "Password set". Example:

```
root PS 2022-05-03 0 99999 7 -1 (Password set, SHA512 crypt.)
```

Remediation:

Set the `root` password with:

```
# passwd root
```

References:

1. NIST SP 800-53 Rev. 5: AC-2, AC-3, AC-5, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1078	TA0005	M1026

5 Logging and Auditing

The items in this section describe how to configure logging, log monitoring, and auditing, using tools included in most distributions.

It is recommended that `rsyslog` be used for logging (with `logwatch` providing summarization) and `auditd` be used for auditing (with `aureport` providing summarization) to automatically monitor logs for intrusion attempts and other suspicious system behavior.

In addition to the local log files created by the steps in this section, it is also recommended that sites collect copies of their system logs on a secure, centralized log server via an encrypted connection. Not only does centralized logging help sites correlate events that may be occurring on multiple systems, but having a second copy of the system log information may be critical after a system compromise where the attacker has modified the local log files on the affected system(s). If a log correlation system is deployed, configure it to process the logs described in this section.

Because it is often necessary to correlate log information from many different systems (particularly after a security incident) it is recommended that the time be synchronized among systems and devices connected to the local network. The standard Internet protocol for time synchronization is the Network Time Protocol (NTP), which is supported by most network-ready devices. Reference <<http://chrony.tuxfamily.org/>> manual page for more information on configuring chrony.

It is important that all logs described in this section be monitored on a regular basis and correlated to determine trends. A seemingly innocuous entry in one log could be more significant when compared to an entry in another log.

Note on log file permissions: There really isn't a "one size fits all" solution to the permissions on log files. Many sites utilize group permissions so that administrators who are in a defined security group, such as "wheel" do not have to elevate privileges to root in order to read log files. Also, if a third party log aggregation tool is used, it may need to have group permissions to read the log files, which is preferable to having it run setuid to root. Therefore, there are two remediation and audit steps for log file permissions. One is for systems that do not have a secured group method implemented that only permits root to read the log files (`root:root 600`). The other is for sites that do have such a setup and are designated as `root:securegrp 640` where `securegrp` is the defined security group (in some cases `wheel`).

5.1 Configure Logging

Logging services should be configured to prevent information leaks and to aggregate logs on a remote server so that they can be reviewed in the event of a system compromise. A centralized log server provides a single point of entry for further analysis, monitoring and filtering.

Security principals for logging

- Ensure transport layer security is implemented between the client and the log server.
- Ensure that logs are rotated as per the environment requirements.
- Ensure all locally generated logs have the appropriate permissions.
- Ensure all security logs are sent to a remote log server.
- Ensure the required events are logged.

What is covered

This section will cover the minimum best practices for the usage of **either** `rsyslog` **or** `journald`. The recommendations are written such that each is wholly independent of each other and **only one is implemented**.

- If your organization makes use of an enterprise wide logging system completely outside of `rsyslog` or `journald`, then the following recommendations do not directly apply. However, the principals of the recommendations should be followed regardless of what solution is implemented. If the enterprise solution incorporates either of these tools, careful consideration should be given to the following recommendations to determine exactly what applies.
- Should your organization make use of both `rsyslog` and `journald`, take care how the recommendations may or may not apply to you.

What is not covered

- Enterprise logging systems not utilizing `rsyslog` or `journald`. As logging is very situational and dependent on the local environment, not everything can be covered here.
- Transport layer security should be applied to all remote logging functionality. Both `rsyslog` and `journald` supports secure transport and should be configured as such.
- The log server. There are a multitude of reasons for a centralized log server (and keeping a short period of logging on the local system), but the log server is out of scope for these recommendations.

5.1.1 Configure rsyslog

The `rsyslog` software package may be used instead of the default `journald` logging mechanism.

Note: This section only applies if `rsyslog` is the chosen method for client side logging. Do not apply this section if `journald` is used.

5.1.1.1 Ensure rsyslog is installed (Manual)

Profile Applicability:

- Level 1 - Server

Description:

The `rsyslog` software is recommended in environments where `journald` does not meet operation requirements.

Rationale:

The security enhancements of `rsyslog` such as connection-oriented (i.e. TCP) transmission of logs, the option to log to database formats, and the encryption of log data en route to a central logging server) justify installing and configuring the package.

Audit:

-IF- rsyslog is being used on the system:

Verify `rsyslog` is installed.

Run the following command:

```
# rpm -q rsyslog
```

Verify the output matches:

```
rsyslog-<version>
```

Remediation:

Run the following command to install `rsyslog`:

```
# dnf install rsyslog
```

References:

1. NIST SP 800-53 Rev. 5: AU-3, AU-12, SI-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			
v7	6.3 <u>Enable Detailed Logging</u> Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1005, T1005.000, T1070, T1070.002	TA0005	

5.1.1.2 Ensure rsyslog service is enabled (Manual)

Profile Applicability:

- Level 1 - Server

Description:

Once the `rsyslog` package is installed, ensure that the service is enabled.

Rationale:

If the `rsyslog` service is not enabled to start on boot, the system will not capture logging events.

Audit:

-IF- rsyslog is being used for logging on the system:
Run the following command to verify `rsyslog` is enabled:

```
# systemctl is-enabled rsyslog
```

Verify the output matches:

```
enabled
```

Remediation:

Run the following command to enable `rsyslog`:

```
# systemctl --now enable rsyslog
```

References:

1. NIST SP 800-53 Rev. 5: AU-3, AU-12

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			
v7	6.3 <u>Enable Detailed Logging</u> Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1211, T1562, T1562.001	TA0005	

5.1.1.3 Ensure journald is configured to send logs to rsyslog (Manual)

Profile Applicability:

- Level 1 - Server

Description:

Data from `journald` may be stored in volatile memory or persisted locally on the server. Utilities exist to accept remote export of `journald` logs, however, use of the `RSyslog` service provides a consistent means of log collection and export.

Rationale:

IF `RSyslog` is the preferred method for capturing logs, all logs of the system should be sent to it for further processing.

Note: This recommendation only applies if `rsyslog` is the chosen method for client side logging. Do not apply this recommendation if `journald` is used.

Audit:

IF `RSyslog` is the preferred method for capturing logs

Review `/etc/systemd/journald.conf` and verify that logs are forwarded to `rsyslog`.

```
# grep ^\s*ForwardToSyslog /etc/systemd/journald.conf
```

Verify the output matches:

```
ForwardToSyslog=yes
```

Remediation:

Edit the `/etc/systemd/journald.conf` file and add the following line:

```
ForwardToSyslog=yes
```

Restart the service:

```
# systemctl restart rsyslog
```

References:

1. NIST SP 800-53 Rev. 5: AC-3, AU-2, AU-4, AU-12, MP-2, SI-5

Additional Information:

As noted in the journald man pages, journald logs may be exported to rsyslog either through the process mentioned here, or through a facility like `systemd-journal.service`. There are trade-offs involved in each implementation, where `ForwardToSyslog` will immediately capture all events (and forward to an external log server, if properly configured), but may not capture all boot-up activities. Mechanisms such as `systemd-journal.service`, on the other hand, will record bootup events, but may delay sending the information to rsyslog, leading to the potential for log manipulation prior to export. Be aware of the limitations of all tools employed to secure a system.

The main configuration file `/etc/systemd/journal.conf` is read before any of the custom `*.conf` files. If there are custom configurations present, they override the main configuration parameters

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.			
v8	8.9 <u>Centralize Audit Logs</u> Centralize, to the extent possible, audit log collection and retention across enterprise assets.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			
v7	6.3 <u>Enable Detailed Logging</u> Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.			
v7	6.5 <u>Central Log Management</u> Ensure that appropriate logs are being aggregated to a central log management system for analysis and review.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1562, T1562.006, T1565	TA0040	M1029

5.1.1.4 Ensure rsyslog default file permissions are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

RSyslog will create logfiles that do not already exist on the system. This setting controls what permissions will be applied to these newly created files.

Rationale:

It is important to ensure that log files have the correct permissions to ensure that sensitive data is archived and protected.

Impact:

The systems global `umask` could override, but only making the file permissions stricter, what is configured in RSyslog with the `FileCreateMode` directive. RSyslog also has its own `$umask` directive that can alter the intended file creation mode. In addition, consideration should be given to how `FileCreateMode` is used.

Thus it is critical to ensure that the intended file creation mode is not overridden with less restrictive settings in `/etc/rsyslog.conf`, `/etc/rsyslog.d/*.conf` files and that `FileCreateMode` is set before any file is created.

Audit:

Run the following command:

```
# grep -Ps '^\\h*\\$FileCreateMode\\h+0[0,2,4,6][0,2,4]0\\b' /etc/rsyslog.conf /etc/rsyslog.d/*.conf
```

Verify the output is includes 0640 or more restrictive:

```
$FileCreateMode 0640
```

Remediation:

Edit either `/etc/rsyslog.conf` or a dedicated `.conf` file in `/etc/rsyslog.d/` and set `$FileCreateMode` to `0640` or more restrictive:

```
$FileCreateMode 0640
```

Restart the service:

```
# systemctl restart rsyslog
```

References:

1. See the `rsyslog.conf(5)` man page for more information.
2. NIST SP 800-53 Rev. 5: AC-3, AC-6, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.			
v7	5.1 <u>Establish Secure Configurations</u> Maintain documented, standard security configuration standards for all authorized operating systems and software.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			
v7	6.3 <u>Enable Detailed Logging</u> Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1083, T1083.000	TA0007	

5.1.1.5 Ensure logging is configured (Manual)

Profile Applicability:

- Level 1 - Server

Description:

The `/etc/rsyslog.conf` and `/etc/rsyslog.d/*.conf` files specifies rules for logging and which files are to be used to log certain classes of messages.

Rationale:

A great deal of important security-related information is sent via `rsyslog` (e.g., successful and failed su attempts, failed login attempts, root login attempts, etc.).

Audit:

Review the contents of `/etc/rsyslog.conf` and `/etc/rsyslog.d/*.conf` files to ensure appropriate logging is set. In addition, run the following command and verify that the log files are logging information as expected:

```
# ls -l /var/log/
```

Remediation:

Edit the following lines in the `/etc/rsyslog.conf` and `/etc/rsyslog.d/*.conf` files as appropriate for your environment.

NOTE: The below configuration is shown for example purposes only. Due care should be given to how the organization wish to store log data.

```
*.emerg                                :omusrmsg:*
auth,authpriv.*                       /var/log/secure
mail.*                                 -/var/log/mail
mail.info                             -/var/log/mail.info
mail.warning                          -/var/log/mail.warn
mail.err                              /var/log/mail.err
cron.*                                /var/log/cron
*.=warning;*.=err                     -/var/log/warn
*.crit                                /var/log/warn
*.*;mail.none;news.none               -/var/log/messages
local0,local1.*                       -/var/log/localmessages
local2,local3.*                       -/var/log/localmessages
local4,local5.*                       -/var/log/localmessages
local6,local7.*                       -/var/log/localmessages
```

Run the following command to reload the `rsyslogd` configuration:

```
# systemctl restart rsyslog
```

References:

1. See the rsyslog.conf(5) man page for more information.
2. NIST SP 800-53 Rev. 5: AU-2, AU-7, AU-12

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			
v7	6.3 <u>Enable Detailed Logging</u> Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002	TA0005	M1047

5.1.1.6 Ensure rsyslog is configured to send logs to a remote log host (Manual)

Profile Applicability:

- Level 1 - Server

Description:

RSyslog supports the ability to send log events it gathers to a remote log host or to receive messages from remote hosts, thus enabling centralized log management.

Rationale:

Storing log data on a remote host protects log integrity from local attacks. If an attacker gains root access on the local system, they could tamper with or remove log data that is stored on the local system.

Audit:

Review the `/etc/rsyslog.conf` and `/etc/rsyslog.d/*.conf` files and verify that logs are sent to a central host (where `loghost.example.com` is the name of your central log host):

Old format

```
# grep "^*.*[^\I][^\I]*@" /etc/rsyslog.conf /etc/rsyslog.d/*.conf
```

Output should include `@@<FQDN or IP of remote loghost>`, for example

```
*.* @@loghost.example.com
```

New format

```
# grep -E '^\\s*([^\#]+\s+)?action\((([^\#]+\s+)?\btarget=\\"?[^\#"]+\\"?\\b' /etc/rsyslog.conf /etc/rsyslog.d/*.conf
```

Output should include `target=<FQDN or IP of remote loghost>`, for example:

```
*.* action(type="omfwd" target="loghost.example.com" port="514" protocol="tcp"
```

Remediation:

Edit the `/etc/rsyslog.conf` and `/etc/rsyslog.d/*.conf` files and add the following line (where `loghost.example.com` is the name of your central log host). The `target` directive may either be a fully qualified domain name or an IP address.

```
*. * action(type="omfwd" target="192.168.2.100" port="514" protocol="tcp"
      action.resumeRetryCount="100"
      queue.type="LinkedList" queue.size="1000")
```

Run the following command to reload the `rsyslogd` configuration:

```
# systemctl restart rsyslog
```

References:

1. See the `rsyslog.conf(5)` man page for more information.
2. NIST SP 800-53 Rev. 5: AU-6

Additional Information:

In addition, see the [RSyslog documentation](#) for implementation details of TLS.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 Collect Audit Logs Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.			
v7	6.2 Activate audit logging Ensure that local logging has been enabled on all systems and networking devices.			
v7	6.3 Enable Detailed Logging Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1562, T1562.006	TA0040	M1029

5.1.1.7 Ensure rsyslog is not configured to receive logs from a remote client (Automated)

Profile Applicability:

- Level 1 - Server

Description:

RSyslog supports the ability to receive messages from remote hosts, thus acting as a log server. Clients should not receive data from other hosts.

Rationale:

If a client is configured to also receive data, thus turning it into a server, the client system is acting outside its operational boundary.

Audit:

Review the `/etc/rsyslog.conf` and `/etc/rsyslog.d/*.conf` files and verify that the system is not configured to accept incoming logs.

New format

```
# grep -Ps -- '^h*module\(load="imtcp"\)' /etc/rsyslog.conf
/etc/rsyslog.d/*.conf
# grep -Ps -- '^h*input\(type="imtcp" port="514"\)' /etc/rsyslog.conf
/etc/rsyslog.d/*.conf
```

No output expected.

-OR-

Old format

```
# grep -s '$ModLoad imtcp' /etc/rsyslog.conf /etc/rsyslog.d/*.conf
# grep -s '$InputTCPServerRun' /etc/rsyslog.conf /etc/rsyslog.d/*.conf
```

No output expected.

Remediation:

Should there be any active log server configuration found in the auditing section, modify those files and remove the specific lines highlighted by the audit. Ensure none of the following entries are present in any of `/etc/rsyslog.conf` or `/etc/rsyslog.d/*.conf`.

New format

```
module(load="imtcp")
input(type="imtcp" port="514")
```

-OR-

Old format

```
$ModLoad imtcp
$InputTCPServerRun
```

Restart the service:

```
# systemctl restart rsyslog
```

References:

1. NIST SP 800-53 Rev. 5: AU-2, AU-7, AU-12, CM-6

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 <u>Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.	●	●	●
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.	●	●	●
v7	6.3 <u>Enable Detailed Logging</u> Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.		●	●
v7	9.2 <u>Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.		●	●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0040	M1029

5.1.2 Configure journald

Included in the systemd suite is a journaling service called `systemd-journald.service` for the collection and storage of logging data. It creates and maintains structured, indexed journals based on logging information that is received from a variety of sources such as:

- Classic RFC3164 BSD syslog via the `/dev/log` socket
- STDOUT/STDERR of programs via `StandardOutput=journal` + `StandardError=journal` in service files (both of which are default settings)
- Kernel log messages via the `/dev/kmsg` device node
- Audit records via the kernel's audit subsystem
- Structured log messages via journald's native protocol

Any changes made to the `systemd-journald` configuration will require a re-start of `systemd-journald`

5.1.2.1 Ensure journald is configured to send logs to a remote log host

5.1.2.1.1 Ensure systemd-journal-remote is installed (Manual)

Profile Applicability:

- Level 1 - Server

Description:

Journald (via `systemd-journal-remote`) supports the ability to send log events it gathers to a remote log host or to receive messages from remote hosts, thus enabling centralized log management.

Rationale:

Storing log data on a remote host protects log integrity from local attacks. If an attacker gains root access on the local system, they could tamper with or remove log data that is stored on the local system.

Audit:

-IF- journald will be used for logging on the system:

Verify `systemd-journal-remote` is installed.

Run the following command:

```
# rpm -q systemd-journal-remote
```

Verify the output matches:

```
systemd-journal-remote-<version>
```

Remediation:

Run the following command to install `systemd-journal-remote`:

```
# dnf install systemd-journal-remote
```

References:

1. NIST SP 800-53 Rev. 5: AU-2, AU-12, SI-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			
v7	6.3 <u>Enable Detailed Logging</u> Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1562, T1562.006	TA0040	M1029

5.1.2.1.2 Ensure systemd-journal-remote is configured (Manual)

Profile Applicability:

- Level 1 - Server

Description:

Journald (via `systemd-journal-remote`) supports the ability to send log events it gathers to a remote log host or to receive messages from remote hosts, thus enabling centralized log management.

Rationale:

Storing log data on a remote host protects log integrity from local attacks. If an attacker gains root access on the local system, they could tamper with or remove log data that is stored on the local system.

Audit:

Verify `systemd-journal-remote` is configured.

Run the following command:

```
# grep -P "^ *URL=|^ *ServerKeyFile=|^ *ServerCertificateFile=|^ *TrustedCertificateFile=" /etc/systemd/journal-upload.conf
```

Verify the output matches per your environments certificate locations and the URL of the log server. Example:

```
URL=192.168.50.42
ServerKeyFile=/etc/ssl/private/journal-upload.pem
ServerCertificateFile=/etc/ssl/certs/journal-upload.pem
TrustedCertificateFile=/etc/ssl/ca/trusted.pem
```

Remediation:

Edit the `/etc/systemd/journal-upload.conf` file and ensure the following lines are set per your environment:

```
URL=192.168.50.42
ServerKeyFile=/etc/ssl/private/journal-upload.pem
ServerCertificateFile=/etc/ssl/certs/journal-upload.pem
TrustedCertificateFile=/etc/ssl/ca/trusted.pem
```

Restart the service:

```
# systemctl restart systemd-journal-upload
```

References:

1. NIST SP 800-53 Rev. 5: AU-2, AU-7 AU-12

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			
v7	6.3 <u>Enable Detailed Logging</u> Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1562, T1562.006	TA0040	M1029

5.1.2.1.3 *Ensure systemd-journal-remote is enabled (Manual)*

Profile Applicability:

- Level 1 - Server

Description:

Journald (via `systemd-journal-remote`) supports the ability to send log events it gathers to a remote log host or to receive messages from remote hosts, thus enabling centralized log management.

Rationale:

Storing log data on a remote host protects log integrity from local attacks. If an attacker gains root access on the local system, they could tamper with or remove log data that is stored on the local system.

Audit:

Verify `systemd-journal-remote` is enabled.

Run the following command:

```
# systemctl is-enabled systemd-journal-upload.service
enabled
```

Remediation:

Run the following command to enable `systemd-journal-remote`:

```
# systemctl --now enable systemd-journal-upload.service
```

References:

1. NIST SP 800-53 Rev. 5: AU-2, AU-12, SI-5, CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			
v7	6.3 <u>Enable Detailed Logging</u> Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1562, T1562.006	TA0040	M1029

5.1.2.1.4 Ensure journald is not configured to receive logs from a remote client (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Journald supports the ability to receive messages from remote hosts, thus acting as a log server. Clients should not receive data from other hosts.

NOTE:

- The same package, `systemd-journal-remote`, is used for both sending logs to remote hosts and receiving incoming logs.
- With regards to receiving logs, there are two services; `systemd-journal-remote.socket` and `systemd-journal-remote.service`.

Rationale:

If a client is configured to also receive data, thus turning it into a server, the client system is acting outside it's operational boundary.

Audit:

Run the following command to verify `systemd-journal-remote.socket` is not enabled:

```
# systemctl is-enabled systemd-journal-remote.socket
```

Verify the output matches:

```
masked
```

Remediation:

Run the following command to disable `systemd-journal-remote.socket`:

```
# systemctl --now mask systemd-journal-remote.socket
```

References:

1. NIST SP 800-53 Rev. 5: AU-2, AU-12, CM-6, CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 <u>Uninstall or Disable Unnecessary Services on Enterprise Assets and Software</u> Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.			
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			
v7	6.3 <u>Enable Detailed Logging</u> Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.			
v7	9.2 <u>Ensure Only Approved Ports, Protocols and Services Are Running</u> Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1562, T1562.006	TA0040	M1029

5.1.2.2 Ensure journald service is enabled (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Ensure that the `systemd-journald` service is enabled to allow capturing of logging events.

Rationale:

If the `systemd-journald` service is not enabled to start on boot, the system will not capture logging events.

Audit:

Run the following command to verify `systemd-journald` is enabled:

```
# systemctl is-enabled systemd-journald.service
```

Verify the output matches:

```
static
```

Remediation:

By default the `systemd-journald` service does not have an `[Install]` section and thus cannot be enabled / disabled. It is meant to be referenced as `Requires` or `Wants` by other unit files. As such, if the status of `systemd-journald` is not `static`, investigate why.

References:

1. NIST SP 800-53 Rev. 5: AU-2, AU-7 AU-12

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			
v7	6.3 <u>Enable Detailed Logging</u> Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1562, T1562.006	TA0040	M1029

5.1.2.3 Ensure journald is configured to compress large log files (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The journald system includes the capability of compressing overly large files to avoid filling up the system with logs or making the logs unmanageably large.

Rationale:

Uncompressed large files may unexpectedly fill a filesystem leading to resource unavailability. Compressing logs prior to write can prevent sudden, unexpected filesystem impacts.

Audit:

Review `/etc/systemd/journald.conf` and verify that large files will be compressed:

```
# grep ^\s*Compress /etc/systemd/journald.conf
```

Verify the output matches:

```
Compress=yes
```

Remediation:

Edit the `/etc/systemd/journald.conf` file and add the following line:

```
Compress=yes
```

Restart the service:

```
# systemctl restart systemd-journald.service
```

References:

1. NIST SP 800-53 Rev. 5: AU-4

Additional Information:

The main configuration file `/etc/systemd/journald.conf` is read before any of the custom `*.conf` files. If there are custom configs present, they override the main configuration parameters.

It is possible to change the default threshold of 512 bytes per object before compression is used.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.			
v8	8.3 <u>Ensure Adequate Audit Log Storage</u> Ensure that logging destinations maintain adequate storage to comply with the enterprise's audit log management process.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			
v7	6.3 <u>Enable Detailed Logging</u> Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.			
v7	6.4 <u>Ensure adequate storage for logs</u> Ensure that all systems that store logs have adequate storage space for the logs generated.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.001	TA0005	M1053

5.1.2.4 Ensure journald is configured to write logfiles to persistent disk (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Data from journald may be stored in volatile memory or persisted locally on the server. Logs in memory will be lost upon a system reboot. By persisting logs to local disk on the server they are protected from loss due to a reboot.

Rationale:

Writing log data to disk will provide the ability to forensically reconstruct events which may have impacted the operations or security of a system even after a system crash or reboot.

Audit:

Review `/etc/systemd/journald.conf` and verify that logs are persisted to disk:

```
# grep ^\s*Storage /etc/systemd/journald.conf
```

Verify the output matches:

```
Storage=persistent
```

Remediation:

Edit the `/etc/systemd/journald.conf` file and add the following line:

```
Storage=persistent
```

Restart the service:

```
# systemctl restart systemd-journald.service
```

References:

1. NIST SP 800-53 Rev. 5: AU-3, AU-12

Additional Information:

The main configuration file `/etc/systemd/journald.conf` is read before any of the custom `*.conf` files. If there are custom configs present, they override the main configuration parameters.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.	●	●	●
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.	●	●	●
v7	6.3 <u>Enable Detailed Logging</u> Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.		●	●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1562, T1562.006	TA0005	M1022

5.1.2.5 Ensure journald is not configured to send logs to rsyslog (Manual)

Profile Applicability:

- Level 1 - Server

Description:

Data from `journald` should be kept in the confines of the service and not forwarded on to other services.

Rationale:

IF `journald` is the method for capturing logs, all logs of the system should be handled by `journald` and not forwarded to other logging mechanisms.

Note: This recommendation only applies if `journald` is the chosen method for client side logging. Do not apply this recommendation if `rsyslog` is used.

Audit:

IF `journald` is the method for capturing logs

Review `/etc/systemd/journald.conf` and verify that logs are not forwarded to `rsyslog`.

```
# grep ^\s*ForwardToSyslog /etc/systemd/journald.conf
```

Verify that there is no output.

Remediation:

Edit the `/etc/systemd/journald.conf` file and ensure that `ForwardToSyslog=yes` is removed.

Restart the service:

```
# systemctl restart systemd-journald.service
```

References:

1. NIST SP 800-53 Rev. 5: AU-2, AU-6, AU-7, AU-12

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.			
v8	8.9 <u>Centralize Audit Logs</u> Centralize, to the extent possible, audit log collection and retention across enterprise assets.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			
v7	6.3 <u>Enable Detailed Logging</u> Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.			
v7	6.5 <u>Central Log Management</u> Ensure that appropriate logs are being aggregated to a central log management system for analysis and review.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1562, T1562.006, T1565	TA0040	M1029

5.1.2.6 Ensure journald log rotation is configured per site policy (Manual)

Profile Applicability:

- Level 1 - Server

Description:

Journald includes the capability of rotating log files regularly to avoid filling up the system with logs or making the logs unmanageably large. The file `/etc/systemd/journald.conf` is the configuration file used to specify how logs generated by Journald should be rotated.

Rationale:

By keeping the log files smaller and more manageable, a system administrator can easily archive these files to another system and spend less time looking through inordinately large log files.

Audit:

Review `/etc/systemd/journald.conf` and verify logs are rotated according to site policy. The specific parameters for log rotation are:

```
SystemMaxUse=  
SystemKeepFree=  
RuntimeMaxUse=  
RuntimeKeepFree=  
MaxFileSec=
```

Remediation:

Review `/etc/systemd/journald.conf` and verify logs are rotated according to site policy. The settings should be carefully understood as there are specific edge cases and prioritization of parameters.

The specific parameters for log rotation are:

```
SystemMaxUse=  
SystemKeepFree=  
RuntimeMaxUse=  
RuntimeKeepFree=  
MaxFileSec=
```

References:

1. NIST SP 800-53 Rev. 5: AU-2, AU-7, AU-12

Additional Information:

See `man 5 journald.conf` for detailed information regarding the parameters in use.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			
v7	6.3 <u>Enable Detailed Logging</u> Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002	TA0040	M1022

5.1.2.7 Ensure journald default file permissions configured (Manual)

Profile Applicability:

- Level 1 - Server

Description:

Journald will create logfiles that do not already exist on the system. This setting controls what permissions will be applied to these newly created files.

Rationale:

It is important to ensure that log files have the correct permissions to ensure that sensitive data is archived and protected.

Audit:

First see if there is an override file `/etc/tmpfiles.d/systemd.conf`. If so, this file will override all default settings as defined in `/usr/lib/tmpfiles.d/systemd.conf` and should be inspected.

If there is no override file, inspect the default `/usr/lib/tmpfiles.d/systemd.conf` against the site specific requirements.

Ensure that file permissions are `0640`.

Should a site policy dictate less restrictive permissions, ensure to follow said policy.

NOTE: More restrictive permissions such as `0600` is implicitly sufficient.

Remediation:

If the default configuration is not appropriate for the site specific requirements, copy `/usr/lib/tmpfiles.d/systemd.conf` to `/etc/tmpfiles.d/systemd.conf` and modify as required. Requirements is either `0640` or site policy if that is less restrictive.

References:

1. NIST SP 800-53 Rev. 5: AC-3, AU-2, AU-12, MP-2, SI-5

Additional Information:

See `man 5 tmpfiles.d` for detailed information on the permission sets for the relevant log files. Further information with examples can be found at

<https://www.freedesktop.org/software/systemd/man/tmpfiles.d.html>

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.			
v7	5.1 <u>Establish Secure Configurations</u> Maintain documented, standard security configuration standards for all authorized operating systems and software.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			
v7	6.3 <u>Enable Detailed Logging</u> Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1083, T1083.000	TA0007	M1022

5.1.3 Ensure all logfiles have appropriate permissions and ownership (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Log files contain information from many services on the local system, or in the event of a centralized log server, others system's logs as well. In general log files are found in `/var/log/`, although application can be configured to store logs elsewhere. Should your application store its logs in another location, ensure to run the same test on that location.

Rationale:

It is important that log files have the correct permissions to ensure that sensitive data is protected and that only the appropriate users / groups have access to them.

Audit:

Run the following script to verify that files in `/var/log/` have appropriate permissions and ownership:

```

#!/usr/bin/env bash

{
    echo -e "\n- Start check - logfiles have appropriate permissions and
ownership"
    output=""
    UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
    find /var/log -type f | (while read -r fname; do
        bname="$(basename "$fname")"
        fugname="$(stat -Lc "%U %G" "$fname")"
        funame="$(awk '{print $1}' <<< "$fugname")"
        fugroup="$(awk '{print $2}' <<< "$fugname")"
        fuid="$(stat -Lc "%u" "$fname")"
        fmode="$(stat -Lc "%a" "$fname")"
        case "$bname" in
            lastlog | lastlog.* | wtmp | wtmp.* | wtmp-* | btmp | btmp.* | btmp-
*)
                if ! grep -Pq -- '^\h*[0,2,4,6][0,2,4,6][0,4]\h*$' <<< "$fmode";
then
                    output="$output\n- File: \"$fname\" mode: \"$fmode\"\n"
                fi
                if ! grep -Pq -- '^\h*root\h+(utmp|root)\h*$' <<< "$fugname";
then
                    output="$output\n- File: \"$fname\" ownership: \"$fugname\"\n"
                fi
                ;;
            secure | auth.log | syslog | messages)
                if ! grep -Pq -- '^\h*[0,2,4,6][0,4]0\h*$' <<< "$fmode"; then
                    output="$output\n- File: \"$fname\" mode: \"$fmode\"\n"
                fi
                if ! grep -Pq -- '^\h*(syslog|root)\h+(adm|root)\h*$' <<<
"$fugname"; then
                    output="$output\n- File: \"$fname\" ownership: \"$fugname\"\n"
                fi
                ;;
            SSSD | sssd)
                if ! grep -Pq -- '^\h*[0,2,4,6][0,2,4,6]0\h*$' <<< "$fmode"; then
                    output="$output\n- File: \"$fname\" mode: \"$fmode\"\n"
                fi
                if ! grep -Piq -- '^\h*(SSSD|root)\h+(SSSD|root)\h*$' <<<
"$fugname"; then
                    output="$output\n- File: \"$fname\" ownership: \"$fugname\"\n"
                fi
                ;;
            gdm | gdm3)
                if ! grep -Pq -- '^\h*[0,2,4,6][0,2,4,6]0\h*$' <<< "$fmode"; then
                    output="$output\n- File: \"$fname\" mode: \"$fmode\"\n"
                fi
                if ! grep -Pq -- '^\h*(root)\h+(gdm3?|root)\h*$' <<< "$fugname";
then
                    output="$output\n- File: \"$fname\" ownership: \"$fugname\"\n"
                fi
                ;;
            *.journal | *.journal~)
                if ! grep -Pq -- '^\h*[0,2,4,6][0,4]0\h*$' <<< "$fmode"; then
                    output="$output\n- File: \"$fname\" mode: \"$fmode\"\n"
                fi
            *)
                ;;
        esac
    done)
}

```

```

        if ! grep -Pq -- '^\\h*(root)\\h+(systemd-journal|root)\\h*$' <<<
"$fugname"; then
            output="$output\\n- File: \"$fname\" ownership: \"$fugname\"\\n"
            fi
        ;;
    *)
        if ! grep -Pq -- '^\\h*[0,2,4,6][0,4]0\\h*$' <<< "$fmode"; then
            output="$output\\n- File: \"$fname\" mode: \"$fmode\"\\n"
            fi
        if [ "$fuid" -ge "$UID_MIN" ] || ! grep -Pq -- '(adm|root|'"$(id
-gn "$funame)"'"')' <<< "$fugroup"; then
            if [ -n "$(awk -v grp="$fugroup" -F: ' $1==grp {print $4}'
/etc/group)" ] || ! grep -Pq '(syslog|root)' <<< "$funame"; then
                output="$output\\n- File: \"$fname\" ownership:
\\$fugname\"\\n"
            fi
        fi
    ;;
esac
done
# If all files passed, then we pass
if [ -z "$output" ]; then
    echo -e "\\n- Audit Results:\\n ** Pass **\\n- All files in \"/var/log/"
have appropriate permissions and ownership\\n"
else
    # print the reason why we are failing
    echo -e "\\n- Audit Results:\\n ** Fail **\\n$output"
fi
echo -e "- End check - logfiles have appropriate permissions and
ownership\\n"
)
}

```

Remediation:

Run the following script to update permissions and ownership on files in `/var/log`. Although the script is not destructive, ensure that the output of the audit procedure is captured in the event that the remediation causes issues.


```
#!/usr/bin/env bash

{
    echo -e "\n- Start remediation - logfiles have appropriate permissions and
ownership"
    UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
    find /var/log -type f | while read -r fname; do
        bname="$(basename "$fname")"
        fugname="$(stat -Lc "%U %G" "$fname")"
        funame="$(awk '{print $1}' <<< "$fugname")"
        fugroup="$(awk '{print $2}' <<< "$fugname")"
        fuid="$(stat -Lc "%u" "$fname")"
        fmode="$(stat -Lc "%a" "$fname")"
        case "$bname" in
            lastlog | lastlog.* | wtmp | wtmp.* | wtmp-* | btmp | btmp.* | btmp-
*)
                ! grep -Pq -- '^\h*[0,2,4,6][0,2,4,6][0,4]\h*$' <<< "$fmode" &&
echo -e "- changing mode on \"$fname\" " && chmod ug-x,o-wx "$fname"
                ! grep -Pq -- '^\h*root\h*$' <<< "$funame" && echo -e "- changing
owner on \"$fname\" " && chown root "$fname"
                ! grep -Pq -- '^\h*(utmp|root)\h*$' <<< "$fugroup" && echo -e "-
changing group on \"$fname\" " && chgrp root "$fname"
                ;;
            secure | auth.log | syslog | messages)
                ! grep -Pq -- '^\h*[0,2,4,6][0,4]0\h*$' <<< "$fmode" && echo -e
"- changing mode on \"$fname\" " && chmod u-x,g-wx,o-rwx "$fname"
                ! grep -Pq -- '^\h*(syslog|root)\h*$' <<< "$funame" && echo -e "-
changing owner on \"$fname\" " && chown root "$fname"
                ! grep -Pq -- '^\h*(adm|root)\h*$' <<< "$fugroup" && echo -e "-
changing group on \"$fname\" " && chgrp root "$fname"
                ;;
            SSSD | sssd)
                ! grep -Pq -- '^\h*[0,2,4,6][0,2,4,6]0\h*$' <<< "$fmode" && echo
-e "- changing mode on \"$fname\" " && chmod ug-x,o-rwx "$fname"
                ! grep -Piq -- '^\h*(SSSD|root)\h*$' <<< "$funame" && echo -e "-
changing owner on \"$fname\" " && chown root "$fname"
                ! grep -Piq -- '^\h*(SSSD|root)\h*$' <<< "$fugroup" && echo -e "-
changing group on \"$fname\" " && chgrp root "$fname"
                ;;
            gdm | gdm3)
                ! grep -Pq -- '^\h*[0,2,4,6][0,2,4,6]0\h*$' <<< "$fmode" && echo
-e "- changing mode on \"$fname\" " && chmod ug-x,o-rwx
                ! grep -Pq -- '^\h*root\h*$' <<< "$funame" && echo -e "- changing
owner on \"$fname\" " && chown root "$fname"
                ! grep -Pq -- '^\h*(gdm3?|root)\h*$' <<< "$fugroup" && echo -e "-
changing group on \"$fname\" " && chgrp root "$fname"
                ;;
            *.journal | *.journal~)
                ! grep -Pq -- '^\h*[0,2,4,6][0,4]0\h*$' <<< "$fmode" && echo -e
"- changing mode on \"$fname\" " && chmod u-x,g-wx,o-rwx "$fname"
                ! grep -Pq -- '^\h*root\h*$' <<< "$funame" && echo -e "- changing
owner on \"$fname\" " && chown root "$fname"
                ! grep -Pq -- '^\h*(systemd-journal|root)\h*$' <<< "$fugroup" &&
echo -e "- changing group on \"$fname\" " && chgrp root "$fname"
                ;;
        *)
            ! grep -Pq -- '^\h*[0,2,4,6][0,4]0\h*$' <<< "$fmode" && echo -e

```

```

"- changing mode on \"$fname\" && chmod u-x,g-wx,o-rwx \"$fname"
    if [ "$fuid" -ge "$UID_MIN" ] || ! grep -Pq -- '(adm|root|'"$(id
-gn "$funame)"')' <<< "$fugroup"; then
        if [ -n "$(awk -v grp="$fugroup" -F: ' $1==grp {print $4}'
/etc/group)" ] || ! grep -Pq '(syslog|root)' <<< "$funame"; then
            [ "$fuid" -ge "$UID_MIN" ] && echo -e "- changing owner on
\"$fname\" && chown root \"$fname"
                ! grep -Pq -- '^\\h*(adm|root)\\h*$' <<< "$fugroup" && echo -
e "- changing group on \"$fname\" && chgrp root \"$fname"
            fi
        fi
    ;;
esac
done
echo -e "- End remediation - logfiles have appropriate permissions and
ownership\n"
}

```

Note: You may also need to change the configuration for your logging software or services for any logs that had incorrect permissions.

If there are services that log to other locations, ensure that those log files have the appropriate permissions.

References:

1. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 Configure Data Access Control Lists Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 Protect Information through Access Control Lists Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1083, T1083.000	TA0007	M1022

5.2 Configure System Accounting (auditd)

The Linux Auditing System operates on a set of rules that collects certain types of system activity to facilitate incident investigation, detect unauthorized access or modification of data. By default events will be logged to `/var/log/audit/audit.log`, which can be configured in `/etc/audit/auditd.conf`.

The following types of audit rules can be specified:

- Control rules: Configuration of the auditing system.
- File system rules: Allow the auditing of access to a particular file or a directory. Also known as file watches.
- System call rules: Allow logging of system calls that any specified program makes.

Audit rules can be set:

- On the command line using the `auditctl` utility. These rules are not persistent across reboots.
- In `/etc/audit/audit.rules`. These rules have to be merged and loaded before they are active.

Notes:

- For 64 bit systems that have `arch` as a rule parameter, you will need two rules: one for 64 bit and one for 32 bit systems calls. For 32 bit systems, only one rule is needed.
- If the auditing system is configured to be locked (`-e 2`), a system reboot will be required in order to load any changes.
- Key names are optional on the rules and will not be used in compliance auditing. The usage of key names is highly recommended as it facilitates organization and searching; as such, all remediation steps will have key names supplied.
- It is best practice to store the rules, in number prepended files, in `/etc/audit/rules.d/`. Rules must end in a `.rules` suffix. This then requires the use of `augenrules` to merge all the rules into `/etc/audit/audit.rules` based on their alphabetical (lexical) sort order. All benchmark recommendations follow this best practice for remediation, specifically using the prefix of `50` which is center weighed if all rule sets make use of the number prepending naming convention.
- Your system may have been customized to change the default `UID_MIN`. All sample output uses `1000`, but this value will not be used in compliance auditing. To confirm the `UID_MIN` for your system, run the following command: `awk '/^\s*UID_MIN/{print $2}' /etc/login.defs`

Normalization

The Audit system normalizes some entries, so when you look at the sample output keep in mind that:

- With regards to users whose login UID is not set, the values `-1 / unset / 4294967295` are equivalent and normalized to `-1`.
- When comparing field types and both sides of the comparison is valid fields types, such as `euid!=uid`, then the auditing system may normalize such that the output is `uid!=euid`.
- Some parts of the rule may be rearranged whilst others are dependent on previous syntax. For example, the following two statements are the same:

```
-a always,exit -F arch=b64 -S execve -C uid!=euid -F auid!=-1 -F  
key=user_emulation
```

and

```
-a always,exit -F arch=b64 -C euid!=uid -F auid!=unset -S execve -k  
user_emulation
```

Capacity planning

The recommendations in this section implement auditing policies that not only produce large quantities of logged data, but may also negatively impact system performance. Capacity planning is critical in order not to adversely impact production environments.

- **Disk space.** If a significantly large set of events are captured, additional on system or off system storage may need to be allocated. If the logs are not sent to a remote log server, ensure that log rotation is implemented else the disk will fill up and the system will halt. Even when logs are sent to a log server, ensure sufficient disk space to allow caching of logs in the case of temporary network outages.
- **Disk IO.** It is not just the amount of data collected that should be considered, but the rate at which logs are generated.
- **CPU overhead.** System call rules might incur considerable CPU overhead. Test the systems open/close syscalls per second with and without the rules to gauge the impact of the rules.

5.2.1 Ensure auditing is enabled

The capturing of system events provides system administrators with information to allow them to determine if unauthorized access to their system is occurring.

5.2.1.1 Ensure auditd is installed (Automated)

Profile Applicability:

- Level 2 - Server

Description:

auditd is the userspace component to the Linux Auditing System. It's responsible for writing audit records to the disk

Rationale:

The capturing of system events provides system administrators with information to allow them to determine if unauthorized access to their system is occurring.

Audit:

Run the following command and verify auditd is installed:

```
# rpm -q audit
```

Remediation:

Run the following command to Install auditd

```
# dnf install audit
```

References:

1. NIST SP 800-53 Rev. 5: AU-2, AU-3, AU-12, SI-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.			
v8	8.5 <u>Collect Detailed Audit Logs</u> Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			
v7	6.3 <u>Enable Detailed Logging</u> Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.001	TA0005	

5.2.1.2 Ensure auditing for processes that start prior to auditd is enabled (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Configure `grub2` so that processes that are capable of being audited can be audited even if they start up prior to `auditd` startup.

Rationale:

Audit events need to be captured on processes that start up prior to `auditd`, so that potential malicious activity cannot go undetected.

Audit:

Run the following command to verify that the `audit=1` parameter has been set:

```
# grubby --info=ALL | grep -Po '\baudit=1\b'
audit=1
```

Note `audit=1` may be returned multiple times

Remediation:

Run the following command to update the `grub2` configuration with `audit=1`:

```
# grubby --update-kernel ALL --args 'audit=1'
```

References:

1. NIST SP 800-53 Rev. 5: AU-2, AU-3, AU-12

Additional Information:

This recommendation is designed around the `grub2` bootloader, if another bootloader is in use in your environment enact equivalent settings.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.001	TA0005	

5.2.1.3 Ensure audit_backlog_limit is sufficient (Automated)

Profile Applicability:

- Level 2 - Server

Description:

The backlog limit has a default setting of 64

Rationale:

During boot if `audit=1`, then the backlog will hold 64 records. If more than 64 records are created during boot, auditd records will be lost and potential malicious activity could go undetected.

Audit:

Run the following command and verify the `audit_backlog_limit=` parameter is set to an appropriate size for your organization

```
# grubby --info=ALL | grep -Po "\baudit_backlog_limit=\d+\b"
audit_backlog_limit=<BACKLOG SIZE>
```

Validate that the line(s) returned contain a value for `audit_backlog_limit=` that is sufficient for your organization.

Recommended that this value be 8192 or larger.

Remediation:

Run the following command to add `audit_backlog_limit=<BACKLOG SIZE>` to `GRUB_CMDLINE_LINUX`:

```
# grubby --update-kernel ALL --args 'audit_backlog_limit=<BACKLOG SIZE>'
```

Example:

```
# grubby --update-kernel ALL --args 'audit_backlog_limit=8192'
```

References:

1. NIST SP 800-53 Rev. 5: AU-2, AU-3, AU-12

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.001	TA0005	

5.2.1.4 Ensure auditd service is enabled (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Turn on the `auditd` daemon to record system events.

Rationale:

The capturing of system events provides system administrators with information to allow them to determine if unauthorized access to their system is occurring.

Audit:

Run the following command to verify `auditd` is enabled:

```
# systemctl is-enabled auditd  
  
enabled
```

Verify result is "enabled".

Remediation:

Run the following command to enable `auditd`:

```
# systemctl --now enable auditd
```

References:

1. NIST SP 800-53 Rev. 5: AU-2, AU-3, AU-12

Additional Information:

Additional methods of enabling a service exist. Consult your distribution documentation for appropriate methods.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			
v7	6.3 <u>Enable Detailed Logging</u> Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.001	TA0005	

5.2.2 Configure Data Retention

When auditing, it is important to carefully configure the storage requirements for audit logs. By default, auditd will max out the log files at 5MB and retain only 4 copies of them. Older versions will be deleted. It is possible on a system that the 20 MBs of audit logs may fill up the system causing loss of audit data. While the recommendations here provide guidance, check your site policy for audit storage requirements.

5.2.2.1 Ensure audit log storage size is configured (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Configure the maximum size of the audit log file. Once the log reaches the maximum size, it will be rotated and a new log file will be started.

Rationale:

It is important that an appropriate size is determined for log files so that they do not impact the system and audit data is not lost.

Audit:

Run the following command and ensure output is in compliance with site policy:

```
# grep -w "^s*max_log_file\s*=" /etc/audit/auditd.conf  
max_log_file = <MB>
```

Remediation:

Set the following parameter in `/etc/audit/auditd.conf` in accordance with site policy:

```
max_log_file = <MB>
```

References:

1. NIST SP 800-53 Rev. 5: AU-8

Additional Information:

The `max_log_file` parameter is measured in megabytes.

Other methods of log rotation may be appropriate based on site policy. One example is time-based rotation strategies which don't have native support in auditd configurations. Manual audit of custom configurations should be evaluated for effectiveness and completeness.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.3 <u>Ensure Adequate Audit Log Storage</u> Ensure that logging destinations maintain adequate storage to comply with the enterprise's audit log management process.			
v7	6.4 <u>Ensure adequate storage for logs</u> Ensure that all systems that store logs have adequate storage space for the logs generated.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0040	M1053

5.2.2.2 Ensure audit logs are not automatically deleted (Automated)

Profile Applicability:

- Level 2 - Server

Description:

The `max_log_file_action` setting determines how to handle the audit log file reaching the max file size. A value of `keep_logs` will rotate the logs but never delete old logs.

Rationale:

In high security contexts, the benefits of maintaining a long audit history exceed the cost of storing the audit history.

Audit:

Run the following command and verify output matches:

```
# grep max_log_file_action /etc/audit/auditd.conf  
  
max_log_file_action = keep_logs
```

Remediation:

Set the following parameter in `/etc/audit/auditd.conf`:

```
max_log_file_action = keep_logs
```

References:

1. NIST SP 800-53 Rev. 5: AU-8

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.3 <u>Ensure Adequate Audit Log Storage</u> Ensure that logging destinations maintain adequate storage to comply with the enterprise's audit log management process.	●	●	●
v7	6.4 <u>Ensure adequate storage for logs</u> Ensure that all systems that store logs have adequate storage space for the logs generated.		●	●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0005	M1028

5.2.2.3 *Ensure system is disabled when audit logs are full* (Automated)

Profile Applicability:

- Level 2 - Server

Description:

The `auditd` daemon can be configured to halt the system when the audit logs are full.

The `admin_space_left_action` parameter tells the system what action to take when the system has detected that it is low on disk space. Valid values are `ignore`, `syslog`, `suspend`, `single`, and `halt`.

- `ignore`, the audit daemon does nothing
- `syslog`, the audit daemon will issue a warning to syslog
- `Suspend`, the audit daemon will stop writing records to the disk
- `single`, the audit daemon will put the computer system in single user mode
- `halt`, the audit daemon will shut down the system

Rationale:

In high security contexts, the risk of detecting unauthorized access or nonrepudiation exceeds the benefit of the system's availability.

Impact:

If the `admin_space_left_action` parameter is set to `halt` the audit daemon will shutdown the system when the disk partition containing the audit logs becomes full.

Audit:

Run the following commands and verify output matches:

```
# grep space_left_action /etc/audit/auditd.conf

space_left_action = email
# grep action_mail_acct /etc/audit/auditd.conf

action_mail_acct = root
```

Run the following command and verify the output is either `halt` **or** `single`:

```
# grep -E 'admin_space_left_action\s*=\s*(halt|single) '
/etc/audit/auditd.conf

admin_space_left_action = <halt|single>
```

Remediation:

Set the following parameters in `/etc/audit/auditd.conf`:

```
space_left_action = email
action_mail_acct = root
```

set `admin_space_left_action` to either `halt` **or** `single` in `/etc/audit/auditd.conf`.

Example:

```
admin_space_left_action = halt
```

References:

1. NIST SP 800-53 Rev. 5: AU-2, AU-8, AU-12, SI-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.	●	●	●
v8	8.3 <u>Ensure Adequate Audit Log Storage</u> Ensure that logging destinations maintain adequate storage to comply with the enterprise's audit log management process.	●	●	●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0005	M1028

5.2.3 Configure auditd rules

The Audit system operates on a set of rules that define what is to be captured in the log files.

The following types of Audit rules can be specified:

- Control rules: Allow the Audit system's behavior and some of its configuration to be modified.
- File system rules: Allow the auditing of access to a particular file or a directory. (Also known as file watches)
- System call rules: Allow logging of system calls that any specified program makes.

Audit rules can be set:

- on the command line using the `auditctl` utility. Note that these rules are not persistent across reboots.
- in a file ending in `.rules` in the `/etc/audit/audit.d/` directory.

5.2.3.1 Ensure changes to system administration scope (sudoers) is collected (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Monitor scope changes for system administrators. If the system has been properly configured to force system administrators to log in as themselves first and then use the `sudo` command to execute privileged commands, it is possible to monitor changes in scope. The file `/etc/sudoers`, or files in `/etc/sudoers.d`, will be written to when the file(s) or related attributes have changed. The audit records will be tagged with the identifier "scope".

Rationale:

Changes in the `/etc/sudoers` and `/etc/sudoers.d` files can indicate that an unauthorized change has been made to the scope of system administrator activity.

Audit:

On disk configuration

Run the following command to check the on disk rules:

```
# awk '/^ *-w/ \
&&/\ /etc\sudoers/ \
&&/ +-p *wa/ \
&&( / key= *[-~]* *$||/ -k *[-~]* *$/) ' /etc/audit/rules.d/*.rules
```

Verify the output matches:

```
-w /etc/sudoers -p wa -k scope
-w /etc/sudoers.d -p wa -k scope
```

Running configuration

Run the following command to check loaded rules:

```
# auditctl -l | awk '/^ *-w/ \
&&/\ /etc\sudoers/ \
&&/ +-p *wa/ \
&&( / key= *[-~]* *$||/ -k *[-~]* *$/) '
```

Verify the output matches:

```
-w /etc/sudoers -p wa -k scope
-w /etc/sudoers.d -p wa -k scope
```

Remediation:

Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor scope changes for system administrators.

Example:

```
# printf "
-w /etc/sudoers -p wa -k scope
-w /etc/sudoers.d -p wa -k scope
" >> /etc/audit/rules.d/50-scope.rules
```

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required.

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. NIST SP 800-53 Rev. 5: AU-3

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 <u>Collect Detailed Audit Logs</u> Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.			
v7	4.8 <u>Log and Alert on Changes to Administrative Group Membership</u> Configure systems to issue a log entry and alert when an account is added to or removed from any group assigned administrative privileges.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0004	M1047

5.2.3.2 Ensure actions as another user are always logged (Automated)

Profile Applicability:

- Level 2 - Server

Description:

`sudo` provides users with temporary elevated privileges to perform operations, either as the superuser or another user.

Rationale:

Creating an audit log of users with temporary elevated privileges and the operation(s) they performed is essential to reporting. Administrators will want to correlate the events written to the audit trail with the records written to `sudo`'s logfile to verify if unauthorized commands have been executed.

Audit:

64 Bit systems

On disk configuration

Run the following command to check the on disk rules:

```
# awk '/^ *-a *always,exit/ \
&&/ -F *arch=b[2346]{2}/ \
&&/ -F *auid!=unset/||/ -F *auid!=-1/||/ -F *auid!=4294967295/) \
&&/ -C *euid!=uid/||/ -C *uid!=euid/) \
&&/ -S *execve/ \
&&/ key= *[^~]* *$/||/ -k *[^~]* *$/)' /etc/audit/rules.d/*.rules
```

Verify the output matches:

```
-a always,exit -F arch=b64 -C euid!=uid -F auid!=unset -S execve -k
user_emulation
-a always,exit -F arch=b32 -C euid!=uid -F auid!=unset -S execve -k
user_emulation
```

Running configuration

Run the following command to check loaded rules:

```
# auditctl -l | awk '/^ *-a *always,exit/ \
&&/ -F *arch=b[2346]{2}/ \
&&/ -F *auid!=unset/||/ -F *auid!=-1/||/ -F *auid!=4294967295/) \
&&/ -C *euid!=uid/||/ -C *uid!=euid/) \
&&/ -S *execve/ \
&&/ key= *[^~]* *$/||/ -k *[^~]* *$/)'
```

Verify the output matches:

```
-a always,exit -F arch=b64 -S execve -C uid!=euid -F auid!=-1 -F
key=user_emulation
-a always,exit -F arch=b32 -S execve -C uid!=euid -F auid!=-1 -F
key=user_emulation
```

32 Bit systems

Follow the same procedures as for 64 bit systems and ignore any entries with b64.

Remediation:

Create audit rules

Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor elevated privileges.

64 Bit systems

Example:

```
# printf "-a always,exit -F arch=b64 -C euid!=uid -F auid!=unset -S execve -k user_emulation\n-a always,exit -F arch=b32 -C euid!=uid -F auid!=unset -S execve -k user_emulation\n" >> /etc/audit/rules.d/50-user_emulation.rules
```

Load audit rules

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required.

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot required to load rules\n"; fi
```

32 Bit systems

Follow the same procedures as for 64 bit systems and ignore any entries with `b64`.

References:

1. NIST SP 800-53 Rev. 5: AU-3

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 <u>Collect Detailed Audit Logs</u> Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.			
v7	4.9 <u>Log and Alert on Unsuccessful Administrative Account Login</u> Configure systems to issue a log entry and alert on unsuccessful logins to an administrative account.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0004	M1047

5.2.3.3 Ensure events that modify the sudo log file are collected (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Monitor the `sudo` log file. If the system has been properly configured to disable the use of the `su` command and force all administrators to have to log in first and then use `sudo` to execute privileged commands, then all administrator commands will be logged to `/var/log/sudo.log`. Any time a command is executed, an audit event will be triggered as the `/var/log/sudo.log` file will be opened for write and the executed administration command will be written to the log.

Rationale:

Changes in `/var/log/sudo.log` indicate that an administrator has executed a command or the log file itself has been tampered with. Administrators will want to correlate the events written to the audit trail with the records written to `/var/log/sudo.log` to verify if unauthorized commands have been executed.

Audit:

On disk configuration

Run the following command to check the on disk rules:

```
# {
  SUDO_LOG_FILE_ESCAPED=$(grep -r logfile /etc/sudoers* | sed -e
  's/.*/logfile=//;s/,? .*/' -e 's"/"/g' -e 's|/|\\|/g')
  [ -n "${SUDO_LOG_FILE_ESCAPED}" ] && awk "/^ *-w/ \
  &&/${SUDO_LOG_FILE_ESCAPED}"/ \
  &&/ +-p *wa/ \
  &&(/ key= *[-~]* *$/||/ -k *[-~]* *$/)" /etc/audit/rules.d/*.rules \
  || printf "ERROR: Variable 'SUDO_LOG_FILE_ESCAPED' is unset.\n"
}
```

Verify output of matches:

```
-w /var/log/sudo.log -p wa -k sudo_log_file
```

Running configuration

Run the following command to check loaded rules:

```
# {
  SUDO_LOG_FILE_ESCAPED=$(grep -r logfile /etc/sudoers* | sed -e
  's/.*/logfile=//;s/,? .*/' -e 's"/"/g' -e 's|/|\\|/g')
  [ -n "${SUDO_LOG_FILE_ESCAPED}" ] && auditctl -l | awk "/^ *-w/ \
  &&/${SUDO_LOG_FILE_ESCAPED}"/ \
  &&/ +-p *wa/ \
  &&(/ key= *[-~]* *$/||/ -k *[-~]* *$/)" \
  || printf "ERROR: Variable 'SUDO_LOG_FILE_ESCAPED' is unset.\n"
}
```

Verify output matches:

```
-w /var/log/sudo.log -p wa -k sudo_log_file
```


Remediation:

Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor events that modify the sudo log file.

Example:

```
# {
SUDO_LOG_FILE=$(grep -r logfile /etc/sudoers* | sed -e 's/.*logfile=//;s/,?
.*//' -e 's/"//g')
[ -n "${SUDO_LOG_FILE}" ] && printf "
-w ${SUDO_LOG_FILE} -p wa -k sudo_log_file
" >> /etc/audit/rules.d/50-sudo.rules || printf "ERROR: Variable
'SUDO_LOG_FILE_ESCAPED' is unset.\n"
}
```

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required.

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. NIST SP 800-53 Rev. 5: AU-3

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 <u>Collect Detailed Audit Logs</u> Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.			
v7	4.9 <u>Log and Alert on Unsuccessful Administrative Account Login</u> Configure systems to issue a log entry and alert on unsuccessful logins to an administrative account.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0004	M1047

5.2.3.4 Ensure events that modify date and time information are collected (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Capture events where the system date and/or time has been modified. The parameters in this section are set to determine if the;

- `adjtimex` - tune kernel clock
- `settimeofday` - set time using `timeval` and `timezone` structures
- `stime` - using seconds since 1/1/1970
- `clock_settime` - allows for the setting of several internal clocks and timers

system calls have been executed. Further, ensure to write an audit record to the configured audit log file upon exit, tagging the records with a unique identifier such as "time-change".

Rationale:

Unexpected changes in system date and/or time could be a sign of malicious activity on the system.

Audit:

64 Bit systems

On disk configuration

Run the following command to check the on disk rules:

```
# {
awk '/^ *-a *always,exit/ \
&&/ -F *arch=b[2346]{2}/ \
&&/ -S/ \
&&/adjtimex/ \
||/settimeofday/ \
||/clock_settime/ ) \
&&/ key= *[-~]* *$/||/ -k *[-~]* *$/)' /etc/audit/rules.d/*.rules

awk '/^ *-w/ \
&&/\etc\localtime/ \
&&/ +-p *wa/ \
&&/ key= *[-~]* *$/||/ -k *[-~]* *$/)' /etc/audit/rules.d/*.rules
}
```

Verify output of matches:

```
-a always,exit -F arch=b64 -S adjtimex,settimeofday,clock_settime -k time-
change
-a always,exit -F arch=b32 -S adjtimex,settimeofday,clock_settime -k time-
change
-w /etc/localtime -p wa -k time-change
```

Running configuration

Run the following command to check loaded rules:

```
# {
auditctl -l | awk '/^ *-a *always,exit/ \
&&/ -F *arch=b[2346]{2}/ \
&&/ -S/ \
&&/adjtimex/ \
||/settimeofday/ \
||/clock_settime/ ) \
&&/ key= *[-~]* *$/||/ -k *[-~]* *$/)'

auditctl -l | awk '/^ *-w/ \
&&/\etc\localtime/ \
&&/ +-p *wa/ \
&&/ key= *[-~]* *$/||/ -k *[-~]* *$/)'
}
```

Verify the output includes:

```
-a always,exit -F arch=b64 -S adjtimex,settimeofday,clock_settime -F
key=time-change
-a always,exit -F arch=b32 -S adjtimex,settimeofday,clock_settime -F
key=time-change
-w /etc/localtime -p wa -k time-change
```

32 Bit systems

Follow the same procedures as for 64 bit systems and ignore any entries with b64. In addition, also audit for the `stime` system call rule. For example:

```
-a always,exit -F arch=b32 -S adjtimex,settimeofday,clock_settime,stime -k time-change
```

Remediation:

Create audit rules

Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor events that modify date and time information.

64 Bit systems

Example:

```
# printf "  
-a always,exit -F arch=b64 -S adjtimex,settimeofday,clock_settime -k time-change  
-a always,exit -F arch=b32 -S adjtimex,settimeofday,clock_settime -k time-change  
-w /etc/localtime -p wa -k time-change  
" >> /etc/audit/rules.d/50-time-change.rules
```

Load audit rules

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required.

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot  
required to load rules\n"; fi
```

32 Bit systems

Follow the same procedures as for 64 bit systems and ignore any entries with b64. In addition, add `stime` to the system call audit. Example:

```
-a always,exit -F arch=b32 -S adjtimex,settimeofday,clock_settime,stime -k time-change
```

References:

1. NIST SP 800-53 Rev. 5: AU-3, CM-6

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 <u>Collect Detailed Audit Logs</u> Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.			
v7	5.5 <u>Implement Automated Configuration Monitoring Systems</u> Utilize a Security Content Automation Protocol (SCAP) compliant configuration monitoring system to verify all security configuration elements, catalog approved exceptions, and alert when unauthorized changes occur.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0005	M1047

5.2.3.5 *Ensure events that modify the system's network environment are collected (Automated)*

Profile Applicability:

- Level 2 - Server

Description:

Record changes to network environment files or system calls. The below parameters monitors the following system calls, and write an audit event on system call exit:

- `sethostname` - set the systems host name
- `setdomainname` - set the systems domain name

The files being monitored are:

- `/etc/issue` and `/etc/issue.net` - messages displayed pre-login
- `/etc/hosts` - file containing host names and associated IP addresses
- `/etc/sysconfig/network` - additional information that is valid to all network interfaces
- `/etc/sysconfig/network-scripts/` - directory containing network interface scripts and configurations files

Rationale:

Monitoring `sethostname` and `setdomainname` will identify potential unauthorized changes to host and domain name of a system. The changing of these names could potentially break security parameters that are set based on those names. The `/etc/hosts` file is monitored for changes that can indicate an unauthorized intruder is trying to change machine associations with IP addresses and trick users and processes into connecting to unintended machines. Monitoring `/etc/issue` and `/etc/issue.net` is important, as intruders could put disinformation into those files and trick users into providing information to the intruder. Monitoring `/etc/sysconfig/network` is important as it can show if network interfaces or scripts are being modified in a way that can lead to the machine becoming unavailable or compromised. All audit records should have a relevant tag associated with them.

Audit:

64 Bit systems

On disk configuration

Run the following commands to check the on disk rules:

```
# {
awk '/^ *-a *always,exit/ \
&&/ -F *arch=b(32|64)/ \
&&/ -S/ \
&&(/sethostname/ \
||/setdomainname/) \
&&(/ key= *![~]* *$/||/ -k *![~]* *$/)' /etc/audit/rules.d/*.rules

awk '/^ *-w/ \
&&(/\etc\issue/ \
||\etc\issue.net/ \
||\etc\hosts/ \
||\etc\sysconfig\network/) \
&&/ +-p *wa/ \
&&(/ key= *![~]* *$/||/ -k *![~]* *$/)' /etc/audit/rules.d/*.rules
}
```

Verify the output matches:

```
-a always,exit -F arch=b64 -S sethostname,setdomainname -k system-locale
-a always,exit -F arch=b32 -S sethostname,setdomainname -k system-locale
-w /etc/issue -p wa -k system-locale
-w /etc/issue.net -p wa -k system-locale
-w /etc/hosts -p wa -k system-locale
-w /etc/sysconfig/network -p wa -k system-locale
-w /etc/sysconfig/network-scripts/ -p wa -k system-locale
```

Running configuration

Run the following command to check loaded rules:

```
# {
auditctl -l | awk '/^ *-a *always,exit/ \
&&/ -F *arch=b(32|64)/ \
&&/ -S/ \
&&(/sethostname/ \
||/setdomainname/) \
&&(/ key= *![~]* *$/||/ -k *![~]* *$/)'

auditctl -l | awk '/^ *-w/ \
&&(/\etc\issue/ \
||\etc\issue.net/ \
||\etc\hosts/ \
||\etc\sysconfig\network/) \
&&/ +-p *wa/ \
&&(/ key= *![~]* *$/||/ -k *![~]* *$/)'
}
```

Verify the output includes:


```
-a always,exit -F arch=b64 -S sethostname,setdomainname -F key=system-locale
-a always,exit -F arch=b32 -S sethostname,setdomainname -F key=system-locale
-w /etc/issue -p wa -k system-locale
-w /etc/issue.net -p wa -k system-locale
-w /etc/hosts -p wa -k system-locale
-w /etc/sysconfig/network -p wa -k system-locale
-w /etc/sysconfig/network-scripts -p wa -k system-locale
```

32 Bit systems

Follow the same procedures as for 64 bit systems and ignore any entries with b64

Remediation:

Create audit rules

Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor events that modify the system's network environment.

64 Bit systems

Example:

```
# printf "
-a always,exit -F arch=b64 -S sethostname,setdomainname -k system-locale
-a always,exit -F arch=b32 -S sethostname,setdomainname -k system-locale
-w /etc/issue -p wa -k system-locale
-w /etc/issue.net -p wa -k system-locale
-w /etc/hosts -p wa -k system-locale
-w /etc/sysconfig/network -p wa -k system-locale
-w /etc/sysconfig/network-scripts/ -p wa -k system-locale
" >> /etc/audit/rules.d/50-system_local.rules
```

Load audit rules

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required.

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

32 Bit systems

Follow the same procedures as for 64 bit systems and ignore any entries with b64.

References:

1. NIST SP 800-53 Rev. 5: AU-3, CM-6

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 <u>Collect Detailed Audit Logs</u> Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.			
v7	5.5 <u>Implement Automated Configuration Monitoring Systems</u> Utilize a Security Content Automation Protocol (SCAP) compliant configuration monitoring system to verify all security configuration elements, catalog approved exceptions, and alert when unauthorized changes occur.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0003	M1047

5.2.3.6 *Ensure use of privileged commands are collected (Automated)*

Profile Applicability:

- Level 2 - Server

Description:

Monitor privileged programs, those that have the `setuid` and/or `setgid` bit set on execution, to determine if unprivileged users are running these commands.

Rationale:

Execution of privileged commands by non-privileged users could be an indication of someone trying to gain unauthorized access to the system.

Impact:

Both the audit and remediation section of this recommendation will traverse all mounted file systems that is not mounted with either `noexec` or `nosuid` mount options. If there are large file systems without these mount options, **such traversal will be significantly detrimental to the performance of the system.**

Before running either the audit or remediation section, inspect the output of the following command to determine exactly which file systems will be traversed:

```
# findmnt -n -l -k -it $(awk '/nodev/ { print $2 }' /proc/filesystems | paste -sd,) | grep -Pv "noexec|nosuid"
```

To exclude a particular file system due to adverse performance impacts, update the audit and remediation sections by adding a sufficiently unique string to the `grep` statement. The above command can be used to test the modified exclusions.

Audit:

On disk configuration

Run the following command to check on disk rules:

```
# for PARTITION in $(findmnt -n -l -k -it $(awk '/nodev/ { print $2 }'
/proc/filesystems | paste -sd,) | grep -Pv "noexec|nosuid" | awk '{print
$1}'); do
    for PRIVILEGED in $(find "${PARTITION}" -xdev -perm /6000 -type f); do
        grep -qr "${PRIVILEGED}" /etc/audit/rules.d && printf "OK:
'${PRIVILEGED}' found in auditing rules.\n" || printf "Warning:
'${PRIVILEGED}' not found in on disk configuration.\n"
    done
done
```

Verify that all output is OK.

Running configuration

Run the following command to check loaded rules:

```
# {
    RUNNING=$(auditctl -l)
    [ -n "${RUNNING}" ] && for PARTITION in $(findmnt -n -l -k -it $(awk
'/nodev/ { print $2 }' /proc/filesystems | paste -sd,) | grep -Pv
"noexec|nosuid" | awk '{print $1}'); do
        for PRIVILEGED in $(find "${PARTITION}" -xdev -perm /6000 -type f); do
            printf -- "${RUNNING}" | grep -q "${PRIVILEGED}" && printf "OK:
'${PRIVILEGED}' found in auditing rules.\n" || printf "Warning:
'${PRIVILEGED}' not found in running configuration.\n"
        done
    done \
    || printf "ERROR: Variable 'RUNNING' is unset.\n"
}
```

Verify that all output is OK.

Special mount points

If there are any special mount points that are not visible by default from `findmnt` as per the above audit, said file systems would have to be manually audited.

Remediation:

Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor the use of privileged commands.

Example:

```
# {
  UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
  AUDIT_RULE_FILE="/etc/audit/rules.d/50-privileged.rules"
  NEW_DATA=()
  for PARTITION in $(findmnt -n -l -k -it $(awk '/nodev/ { print $2 }'
/proc/filesystems | paste -sd,) | grep -Pv "noexec|nosuid" | awk '{print
$1}'); do
    readarray -t DATA <<(find "${PARTITION}" -xdev -perm /6000 -type f | awk
-v UID_MIN=${UID_MIN} '{print "-a always,exit -F path=" $1 " -F perm=x -F
audit>="UID_MIN" -F audit!=unset -k privileged" }')
    for ENTRY in "${DATA[@]"}"; do
      NEW_DATA+=("${ENTRY}")
    done
  done
  readarray &> /dev/null -t OLD_DATA < "${AUDIT_RULE_FILE}"
  COMBINED_DATA=( "${OLD_DATA[@]}" "${NEW_DATA[@]}" )
  printf '%s\n' "${COMBINED_DATA[@]}" | sort -u > "${AUDIT_RULE_FILE}"
}
```

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required.

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

Special mount points

If there are any special mount points that are not visible by default from just scanning `/`, change the `PARTITION` variable to the appropriate partition and re-run the remediation.

References:

1. NIST SP 800-53 Rev. 5: AU-3

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 <u>Collect Detailed Audit Logs</u> Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0002	M1026

5.2.3.7 Ensure unsuccessful file access attempts are collected (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Monitor for unsuccessful attempts to access files. The following parameters are associated with system calls that control files:

- **creation** - `creat`
- **opening** - `open` , `openat`
- **truncation** - `truncate` , `ftruncate`

An audit log record will only be written if all of the following criteria is met for the user when trying to access a file:

- a non-privileged user (`auid >= UID_MIN`)
- is not a Daemon event (`auid = 4294967295/unset/-1`)
- if the system call returned `EACCES` (permission denied) or `EPERM` (some other permanent error associated with the specific system call)

Rationale:

Failed attempts to open, create or truncate files could be an indication that an individual or process is trying to gain unauthorized access to the system.

Audit:

64 Bit systems

On disk configuration

Run the following command to check the on disk rules:

```
# {
  UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
  [ -n "${UID_MIN}" ] && awk "/^ *-a *always,exit/ \
&&/ -F *arch=b[2346]{2}/ \
&&(/ -F *aud!=unset/||/ -F *aud!=-1/||/ -F *aud!=4294967295/) \
&&/ -F *aud>=${UID_MIN}/ \
&&(/ -F *exit==EACCES/||/ -F *exit==EPERM/) \
&&/ -S/ \
&&/creat/ \
&&/open/ \
&&/truncate/ \
&&(/ key= *[:-~]* *$/||/ -k *[:-~]* *$/)" /etc/audit/rules.d/*.rules \
  || printf "ERROR: Variable 'UID_MIN' is unset.\n"
}
```

Verify the output includes:

```
-a always,exit -F arch=b64 -S creat,open,openat,truncate,ftruncate -F exit==
EACCES -F aud>=1000 -F aud!=unset -k access
-a always,exit -F arch=b64 -S creat,open,openat,truncate,ftruncate -F exit==
EPERM -F aud>=1000 -F aud!=unset -k access
-a always,exit -F arch=b32 -S creat,open,openat,truncate,ftruncate -F exit==
EACCES -F aud>=1000 -F aud!=unset -k access
-a always,exit -F arch=b32 -S creat,open,openat,truncate,ftruncate -F exit==
EPERM -F aud>=1000 -F aud!=unset -k access
```

Running configuration

Run the following command to check loaded rules:

```
# {
  UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
  [ -n "${UID_MIN}" ] && auditctl -l | awk "/^ *-a *always,exit/ \
&&/ -F *arch=b[2346]{2}/ \
&&(/ -F *aud!=unset/||/ -F *aud!=-1/||/ -F *aud!=4294967295/) \
&&/ -F *aud>=${UID_MIN}/ \
&&(/ -F *exit==EACCES/||/ -F *exit==EPERM/) \
&&/ -S/ \
&&/creat/ \
&&/open/ \
&&/truncate/ \
&&(/ key= *[:-~]* *$/||/ -k *[:-~]* *$/)" \
  || printf "ERROR: Variable 'UID_MIN' is unset.\n"
}
```

Verify the output includes:


```
-a always,exit -F arch=b64 -S open,truncate,ftruncate,creat,openat -F exit=-EACCES -F auid>=1000 -F auid!=-1 -F key=access
-a always,exit -F arch=b64 -S open,truncate,ftruncate,creat,openat -F exit=-EPERM -F auid>=1000 -F auid!=-1 -F key=access
-a always,exit -F arch=b32 -S open,truncate,ftruncate,creat,openat -F exit=-EACCES -F auid>=1000 -F auid!=-1 -F key=access
-a always,exit -F arch=b32 -S open,truncate,ftruncate,creat,openat -F exit=-EPERM -F auid>=1000 -F auid!=-1 -F key=access
```

32 Bit systems

Follow the same procedures as for 64 bit systems and ignore any entries with b64.

Remediation:

Create audit rules

Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor unsuccessful file access attempts.

64 Bit systems

Example:

```
# {
UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
[ -n "${UID_MIN}" ] && printf "
-a always,exit -F arch=b64 -S creat,open,openat,truncate,ftruncate -F exit=-EACCES -F auid>=${UID_MIN} -F auid!=unset -k access
-a always,exit -F arch=b64 -S creat,open,openat,truncate,ftruncate -F exit=-EPERM -F auid>=${UID_MIN} -F auid!=unset -k access
-a always,exit -F arch=b32 -S creat,open,openat,truncate,ftruncate -F exit=-EACCES -F auid>=${UID_MIN} -F auid!=unset -k access
-a always,exit -F arch=b32 -S creat,open,openat,truncate,ftruncate -F exit=-EPERM -F auid>=${UID_MIN} -F auid!=unset -k access
" >> /etc/audit/rules.d/50-access.rules || printf "ERROR: Variable 'UID_MIN' is unset.\n"
}
```

Load audit rules

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required.

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot required to load rules\n"; fi
```

32 Bit systems

Follow the same procedures as for 64 bit systems and ignore any entries with b64.

References:

1. NIST SP 800-53 Rev. 5: AU-3

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 Collect Detailed Audit Logs Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.			
v7	14.9 Enforce Detail Logging for Access or Changes to Sensitive Data Enforce detailed audit logging for access to sensitive data or changes to sensitive data (utilizing tools such as File Integrity Monitoring or Security Information and Event Monitoring).			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0007	M1047

5.2.3.8 Ensure events that modify user/group information are collected (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Record events affecting the modification of user or group information, including that of passwords and old passwords if in use.

- `/etc/group` - system groups
- `/etc/passwd` - system users
- `/etc/gshadow` - encrypted password for each group
- `/etc/shadow` - system user passwords
- `/etc/security/opasswd` - storage of old passwords if the relevant PAM module is in use

The parameters in this section will watch the files to see if they have been opened for write or have had attribute changes (e.g. permissions) and tag them with the identifier "identity" in the audit log file.

Rationale:

Unexpected changes to these files could be an indication that the system has been compromised and that an unauthorized user is attempting to hide their activities or compromise additional accounts.

Audit:

On disk configuration

Run the following command to check the on disk rules:

```
# awk '/^ *-w/ \
&&(/\/etc\/group/ \
  ||\/etc\/passwd/ \
  ||\/etc\/gshadow/ \
  ||\/etc\/shadow/ \
  ||\/etc\/security\/opasswd/) \
&&/ +-p *wa/ \
&&(/ key= *[-~]* *$/||/ -k *[-~]* *$/)' /etc/audit/rules.d/*.rules
```

Verify the output matches:

```
-w /etc/group -p wa -k identity
-w /etc/passwd -p wa -k identity
-w /etc/gshadow -p wa -k identity
-w /etc/shadow -p wa -k identity
-w /etc/security/opasswd -p wa -k identity
```

Running configuration

Run the following command to check loaded rules:

```
# auditctl -l | awk '/^ *-w/ \
&&(/\/etc\/group/ \
  ||\/etc\/passwd/ \
  ||\/etc\/gshadow/ \
  ||\/etc\/shadow/ \
  ||\/etc\/security\/opasswd/) \
&&/ +-p *wa/ \
&&(/ key= *[-~]* *$/||/ -k *[-~]* *$/)'
```

Verify the output matches:

```
-w /etc/group -p wa -k identity
-w /etc/passwd -p wa -k identity
-w /etc/gshadow -p wa -k identity
-w /etc/shadow -p wa -k identity
-w /etc/security/opasswd -p wa -k identity
```

Remediation:

Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor events that modify user/group information.

Example:

```
# printf "  
-w /etc/group -p wa -k identity  
-w /etc/passwd -p wa -k identity  
-w /etc/gshadow -p wa -k identity  
-w /etc/shadow -p wa -k identity  
-w /etc/security/opasswd -p wa -k identity  
" >> /etc/audit/rules.d/50-identity.rules
```

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required.

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot  
required to load rules\n"; fi
```

References:

1. NIST SP 800-53 Rev. 5: AU-3

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 <u>Collect Detailed Audit Logs</u> Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.			
v7	4.8 <u>Log and Alert on Changes to Administrative Group Membership</u> Configure systems to issue a log entry and alert when an account is added to or removed from any group assigned administrative privileges.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0004	M1047

5.2.3.9 Ensure discretionary access control permission modification events are collected (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Monitor changes to file permissions, attributes, ownership and group. The parameters in this section track changes for system calls that affect file permissions and attributes. The following commands and system calls effect the permissions, ownership and various attributes of files.

- `chmod`
- `fchmod`
- `fchmodat`
- `chown`
- `fchown`
- `fchownat`
- `lchown`
- `setxattr`
- `lsetxattr`
- `fsetxattr`
- `removexattr`
- `lremovexattr`
- `fremovexattr`

In all cases, an audit record will only be written for non-system user ids and will ignore Daemon events. All audit records will be tagged with the identifier "perm_mod."

Rationale:

Monitoring for changes in file attributes could alert a system administrator to activity that could indicate intruder activity or policy violation.

Audit:

64 Bit systems

On disk configuration

Run the following command to check the on disk rules:

```
# {
  UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
  [ -n "${UID_MIN}" ] && awk "/^ *-a *always,exit/ \
&&/ -F *arch=b[2346]{2}/ \
&&/ -F *auid!=unset/||/ -F *auid!=-1/||/ -F *auid!=4294967295/) \
&&/ -S/ \
&&/ -F *auid>=${UID_MIN}/ \
&&(/chmod/||/fchmod/||/fchmodat/ \
  ||/chown/||/fchown/||/fchownat/||/lchown/ \
  ||/setxattr/||/lsetxattr/||/fsetxattr/ \
  ||/removexattr/||/lremovexattr/||/fremovexattr/) \
&&/ key= *[:-~]* *$/||/ -k *[:-~]* *$/)" /etc/audit/rules.d/*.rules \
  || printf "ERROR: Variable 'UID_MIN' is unset.\n"
}
```

Verify the output matches:

```
-a always,exit -F arch=b64 -S chmod,fchmod,fchmodat -F auid>=1000 -F
auid!=unset -F key=perm_mod
-a always,exit -F arch=b64 -S chown,fchown,lchown,fchownat -F auid>=1000 -F
auid!=unset -F key=perm_mod
-a always,exit -F arch=b32 -S chmod,fchmod,fchmodat -F auid>=1000 -F
auid!=unset -F key=perm_mod
-a always,exit -F arch=b32 -S lchown,fchown,chown,fchownat -F auid>=1000 -F
auid!=unset -F key=perm_mod
-a always,exit -F arch=b64 -S
setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr -F
auid>=1000 -F auid!=unset -F key=perm_mod
-a always,exit -F arch=b32 -S
setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr -F
auid>=1000 -F auid!=unset -F key=perm_mod
```

Running configuration

Run the following command to check loaded rules:


```
# {
UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
[ -n "${UID_MIN}" ] && auditctl -l | awk "/^ *-a *always,exit/ \
&&/ -F *arch=b[2346]{2}/ \
&&/ -F *auid!=unset/||/ -F *auid!=-1/||/ -F *auid!=4294967295/) \
&&/ -S/ \
&&/ -F *auid>=${UID_MIN}/ \
&&(/chmod/||/fchmod/||/fchmodat/ \
||/chown/||/fchown/||/fchownat/||/lchown/ \
||/setxattr/||/lsetxattr/||/fsetxattr/ \
||/removexattr/||/lremovexattr/||/fremovexattr/) \
&&/ key= *[:-~]* *$/||/ -k *[:-~]* *$/)" \
|| printf "ERROR: Variable 'UID_MIN' is unset.\n"
}
```

Verify the output matches:

```
-a always,exit -F arch=b64 -S chmod,fchmod,fchmodat -F auid>=1000 -F auid!=-1
-F key=perm_mod
-a always,exit -F arch=b64 -S chown,fchown,lchown,fchownat -F auid>=1000 -F
auid!=-1 -F key=perm_mod
-a always,exit -F arch=b32 -S chmod,fchmod,fchmodat -F auid>=1000 -F auid!=-1
-F key=perm_mod
-a always,exit -F arch=b32 -S lchown,fchown,chown,fchownat -F auid>=1000 -F
auid!=-1 -F key=perm_mod
-a always,exit -F arch=b64 -S
setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr -F
auid>=1000 -F auid!=-1 -F key=perm_mod
-a always,exit -F arch=b32 -S
setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr -F
auid>=1000 -F auid!=-1 -F key=perm_mod
```

32 Bit systems

Follow the same procedures as for 64 bit systems and ignore any entries with b64.

Remediation:

Create audit rules

Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor discretionary access control permission modification events.

64 Bit systems

Example:

```
# {
UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
[ -n "${UID_MIN}" ] && printf "
-a always,exit -F arch=b64 -S chmod,fchmod,fchmodat -F auid>=${UID_MIN} -F
auid!=unset -F key=perm_mod
-a always,exit -F arch=b64 -S chown,fchown,lchown,fchownat -F
auid>=${UID_MIN} -F auid!=unset -F key=perm_mod
-a always,exit -F arch=b32 -S chmod,fchmod,fchmodat -F auid>=${UID_MIN} -F
auid!=unset -F key=perm_mod
-a always,exit -F arch=b32 -S lchown,fchown,chown,fchownat -F
auid>=${UID_MIN} -F auid!=unset -F key=perm_mod
-a always,exit -F arch=b64 -S
setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr -F
auid>=${UID_MIN} -F auid!=unset -F key=perm_mod
-a always,exit -F arch=b32 -S
setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr -F
auid>=${UID_MIN} -F auid!=unset -F key=perm_mod
" >> /etc/audit/rules.d/50-perm_mod.rules || printf "ERROR: Variable
'UID_MIN' is unset.\n"
}
```

Load audit rules

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required.

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

32 Bit systems

Follow the same procedures as for 64 bit systems and ignore any entries with `b64`.

References:

1. NIST SP 800-53 Rev. 5: AU-3, CM-6

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 <u>Collect Detailed Audit Logs</u> Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.			
v7	5.5 <u>Implement Automated Configuration Monitoring Systems</u> Utilize a Security Content Automation Protocol (SCAP) compliant configuration monitoring system to verify all security configuration elements, catalog approved exceptions, and alert when unauthorized changes occur.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0005	M1022

5.2.3.10 Ensure successful file system mounts are collected (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Monitor the use of the `mount` system call. The `mount` (and `umount`) system call controls the mounting and unmounting of file systems. The parameters below configure the system to create an audit record when the mount system call is used by a non-privileged user

Rationale:

It is highly unusual for a non privileged user to `mount` file systems to the system. While tracking `mount` commands gives the system administrator evidence that external media may have been mounted (based on a review of the source of the mount and confirming it's an external media type), it does not conclusively indicate that data was exported to the media. System administrators who wish to determine if data were exported, would also have to track successful `open`, `creat` and `truncate` system calls requiring write access to a file under the mount point of the external media file system. This could give a fair indication that a write occurred. The only way to truly prove it, would be to track successful writes to the external media. Tracking write system calls could quickly fill up the audit log and is not recommended. Recommendations on configuration options to track data export to media is beyond the scope of this document.

Audit:

64 Bit systems

On disk configuration

Run the following command to check the on disk rules:

```
# {
  UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
  [ -n "${UID_MIN}" ] && awk "/^ *-a *always,exit/ \
&&/ -F *arch=b[2346]{2}/ \
&&(/ -F *auid!=unset/||/ -F *auid!=-1/||/ -F *auid!=4294967295/) \
&&/ -F *auid>=${UID_MIN}/ \
&&/ -S/ \
&&/mount/ \
&&(/ key= *[-~]* *$/||/ -k *[-~]* *$/)" /etc/audit/rules.d/*.rules \
  || printf "ERROR: Variable 'UID_MIN' is unset.\n"
}
```

Verify the output matches:

```
-a always,exit -F arch=b64 -S mount -F auid>=1000 -F auid!=unset -k mounts
-a always,exit -F arch=b32 -S mount -F auid>=1000 -F auid!=unset -k mounts
```

Running configuration

Run the following command to check loaded rules:

```
# {
  UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
  [ -n "${UID_MIN}" ] && auditctl -l | awk "/^ *-a *always,exit/ \
&&/ -F *arch=b[2346]{2}/ \
&&(/ -F *auid!=unset/||/ -F *auid!=-1/||/ -F *auid!=4294967295/) \
&&/ -F *auid>=${UID_MIN}/ \
&&/ -S/ \
&&/mount/ \
&&(/ key= *[-~]* *$/||/ -k *[-~]* *$/)" \
  || printf "ERROR: Variable 'UID_MIN' is unset.\n"
}
```

Verify the output matches:

```
-a always,exit -F arch=b64 -S mount -F auid>=1000 -F auid!=-1 -F key=mounts
-a always,exit -F arch=b32 -S mount -F auid>=1000 -F auid!=-1 -F key=mounts
```

32 Bit systems

Follow the same procedures as for 64 bit systems and ignore any entries with b64.

Remediation:

Create audit rules

Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor successful file system mounts.

64 Bit systems

Example:

```
# {
UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
[ -n "${UID_MIN}" ] && printf "
-a always,exit -F arch=b32 -S mount -F auid>=1000 -F auid!=unset -k mounts
-a always,exit -F arch=b64 -S mount -F auid>=1000 -F auid!=unset -k mounts
" >> /etc/audit/rules.d/50-mounts.rules || printf "ERROR: Variable 'UID_MIN'
is unset.\n"
}
```

Load audit rules

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required.

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

32 Bit systems

Follow the same procedures as for 64 bit systems and ignore any entries with `b64`.

References:

1. NIST SP 800-53 Rev. 5: CM-6

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 <u>Collect Detailed Audit Logs</u> Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.			
v7	6.3 <u>Enable Detailed Logging</u> Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0010	M1034

5.2.3.11 Ensure session initiation information is collected (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Monitor session initiation events. The parameters in this section track changes to the files associated with session events.

- `/var/run/utmp` - tracks all currently logged in users.
- `/var/log/wtmp` - file tracks logins, logouts, shutdown, and reboot events.
- `/var/log/btmp` - keeps track of failed login attempts and can be read by entering the command `/usr/bin/last -f /var/log/btmp`.

All audit records will be tagged with the identifier "session."

Rationale:

Monitoring these files for changes could alert a system administrator to logins occurring at unusual hours, which could indicate intruder activity (i.e. a user logging in at a time when they do not normally log in).

Audit:

On disk configuration

Run the following command to check the on disk rules:

```
# awk '/^ *-w/ \
&&(/\/var\/run\/utmp/ \
  ||\/var\/log\/wtmp/ \
  ||\/var\/log\/btmp/) \
&&/ +-p *wa/ \
&&(/ key= *![~]* *$/||/ -k *![~]* *$/)' /etc/audit/rules.d/*.rules
```

Verify the output matches:

```
-w /var/run/utmp -p wa -k session
-w /var/log/wtmp -p wa -k session
-w /var/log/btmp -p wa -k session
```

Running configuration

Run the following command to check loaded rules:

```
# auditctl -l | awk '/^ *-w/ \
&&(/\/var\/run\/utmp/ \
  ||\/var\/log\/wtmp/ \
  ||\/var\/log\/btmp/) \
&&/ +-p *wa/ \
&&(/ key= *![~]* *$/||/ -k *![~]* *$/)'
```

Verify the output matches:

```
-w /var/run/utmp -p wa -k session
-w /var/log/wtmp -p wa -k session
-w /var/log/btmp -p wa -k session
```

Remediation:

Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor session initiation information.

Example:

```
# printf "  
-w /var/run/utmp -p wa -k session  
-w /var/log/wtmp -p wa -k session  
-w /var/log/btmp -p wa -k session  
" >> /etc/audit/rules.d/50-session.rules
```

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required.

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot  
required to load rules\n"; fi
```

References:

1. NIST SP 800-53 Rev. 5: AU-3

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 <u>Collect Detailed Audit Logs</u> Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.			
v7	4.9 <u>Log and Alert on Unsuccessful Administrative Account Login</u> Configure systems to issue a log entry and alert on unsuccessful logins to an administrative account.			
v7	16.13 <u>Alert on Account Login Behavior Deviation</u> Alert when users deviate from normal login behavior, such as time-of-day, workstation location and duration.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0001	M1047

5.2.3.12 Ensure login and logout events are collected (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Monitor login and logout events. The parameters below track changes to files associated with login/logout events.

- `/var/log/lastlog` - maintain records of the last time a user successfully logged in.
- `/var/run/faillock` - directory maintains records of login failures via the `pam_faillock` module.

Rationale:

Monitoring login/logout events could provide a system administrator with information associated with brute force attacks against user logins.

Audit:

On disk configuration

Run the following command to check the on disk rules:

```
# awk '/^ *-w/ \
&&(/\/var\/log\/lastlog/ \
  ||\/var\/run\/faillock/) \
&&/ +-p *wa/ \
&&(/ key= *![~]* *$/||/ -k *![~]* *$/)' /etc/audit/rules.d/*.rules
```

Verify the output matches:

```
-w /var/log/lastlog -p wa -k logins
-w /var/run/faillock -p wa -k logins
```

Running configuration

Run the following command to check loaded rules:

```
# auditctl -l | awk '/^ *-w/ \
&&(/\/var\/log\/lastlog/ \
  ||\/var\/run\/faillock/) \
&&/ +-p *wa/ \
&&(/ key= *![~]* *$/||/ -k *![~]* *$/)'
```

Verify the output matches:

```
-w /var/log/lastlog -p wa -k logins
-w /var/run/faillock -p wa -k logins
```

Remediation:

Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor login and logout events.

Example:

```
# printf "
-w /var/log/lastlog -p wa -k logins
-w /var/run/faillock -p wa -k logins
" >> /etc/audit/rules.d/50-login.rules
```

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required.

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. NIST SP 800-53 Rev. 5: AU-3

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 <u>Collect Detailed Audit Logs</u> Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.			
v7	4.9 <u>Log and Alert on Unsuccessful Administrative Account Login</u> Configure systems to issue a log entry and alert on unsuccessful logins to an administrative account.			
v7	16.11 <u>Lock Workstation Sessions After Inactivity</u> Automatically lock workstation sessions after a standard period of inactivity.			
v7	16.13 <u>Alert on Account Login Behavior Deviation</u> Alert when users deviate from normal login behavior, such as time-of-day, workstation location and duration.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0001	M1047

5.2.3.13 Ensure file deletion events by users are collected (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Monitor the use of system calls associated with the deletion or renaming of files and file attributes. This configuration statement sets up monitoring for:

- `unlink` - remove a file
- `unlinkat` - remove a file attribute
- `rename` - rename a file
- `renameat` rename a file attribute system calls and tags them with the identifier "delete".

Rationale:

Monitoring these calls from non-privileged users could provide a system administrator with evidence that inappropriate removal of files and file attributes associated with protected files is occurring. While this audit option will look at all events, system administrators will want to look for specific privileged files that are being deleted or altered.

Audit:

64 Bit systems

On disk configuration

Run the following command to check the on disk rules:

```
# {
  UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
  [ -n "${UID_MIN}" ] && awk "/^ *-a *always,exit/ \
    &&/ -F *arch=b[2346]{2}/ \
    &&(/ -F *auid!=unset/||/ -F *auid!=-1/||/ -F *auid!=4294967295/) \
    &&/ -F *auid>=${UID_MIN})/ \
    &&/ -S/ \
    &&(/unlink/||/rename/||/unlinkat/||/renameat/) \
    &&(/ key= *[-~]* *$/||/ -k *[-~]* *$/)" /etc/audit/rules.d/*.rules \
    || printf "ERROR: Variable 'UID_MIN' is unset.\n"
}
```

Verify the output matches:

```
-a always,exit -F arch=b64 -S unlink,unlinkat,rename,renameat -F auid>=1000 -
F auid!=unset -k delete
-a always,exit -F arch=b32 -S unlink,unlinkat,rename,renameat -F auid>=1000 -
F auid!=unset -k delete
```

Running configuration

Run the following command to check loaded rules:

```
# {
  UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
  [ -n "${UID_MIN}" ] && auditctl -l | awk "/^ *-a *always,exit/ \
    &&/ -F *arch=b[2346]{2}/ \
    &&(/ -F *auid!=unset/||/ -F *auid!=-1/||/ -F *auid!=4294967295/) \
    &&/ -F *auid>=${UID_MIN})/ \
    &&/ -S/ \
    &&(/unlink/||/rename/||/unlinkat/||/renameat/) \
    &&(/ key= *[-~]* *$/||/ -k *[-~]* *$/)" \
    || printf "ERROR: Variable 'UID_MIN' is unset.\n"
}
```

Verify the output matches:

```
-a always,exit -F arch=b64 -S rename,unlink,unlinkat,renameat -F auid>=1000 -
F auid!=-1 -F key=delete
-a always,exit -F arch=b32 -S unlink,rename,unlinkat,renameat -F auid>=1000 -
F auid!=-1 -F key=delete
```

32 Bit systems

Follow the same procedures as for 64 bit systems and ignore any entries with b64.

Remediation:

Create audit rules

Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor file deletion events by users.

64 Bit systems

Example:

```
# {
UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
[ -n "${UID_MIN}" ] && printf "
-a always,exit -F arch=b64 -S rename,unlink,unlinkat,renameat -F
auid>=${UID_MIN} -F auid!=unset -F key=delete
-a always,exit -F arch=b32 -S rename,unlink,unlinkat,renameat -F
auid>=${UID_MIN} -F auid!=unset -F key=delete
" >> /etc/audit/rules.d/50-delete.rules || printf "ERROR: Variable 'UID_MIN'
is unset.\n"
}
```

Load audit rules

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required.

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

32 Bit systems

Follow the same procedures as for 64 bit systems and ignore any entries with `b64`.

References:

1. NIST SP 800-53 Rev. 5: AU-12, SC-7

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 <u>Collect Detailed Audit Logs</u> Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0005	M1047

5.2.3.14 Ensure events that modify the system's Mandatory Access Controls are collected (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Monitor SELinux, an implementation of mandatory access controls. The parameters below monitor any write access (potential additional, deletion or modification of files in the directory) or attribute changes to the `/etc/selinux/` and `/usr/share/selinux/` directories.

Note: If a different Mandatory Access Control method is used, changes to the corresponding directories should be audited.

Rationale:

Changes to files in the `/etc/selinux/` and `/usr/share/selinux/` directories could indicate that an unauthorized user is attempting to modify access controls and change security contexts, leading to a compromise of the system.

Audit:

On disk configuration

Run the following command to check the on disk rules:

```
# awk '/^ *-w/ \
&&(/\/etc\/selinux/ \
  ||\/usr\/share\/selinux/) \
&&/ +-p *wa/ \
&&(/ key= *![~]* *$/||/ -k *![~]* *$/)' /etc/audit/rules.d/*.rules
```

Verify the output matches:

```
-w /etc/selinux -p wa -k MAC-policy
-w /usr/share/selinux -p wa -k MAC-policy
```

Running configuration

Run the following command to check loaded rules:

```
# auditctl -l | awk '/^ *-w/ \
&&(/\/etc\/selinux/ \
  ||\/usr\/share\/selinux/) \
&&/ +-p *wa/ \
&&(/ key= *![~]* *$/||/ -k *![~]* *$/)'
```

Verify the output matches:

```
-w /etc/selinux -p wa -k MAC-policy
-w /usr/share/selinux -p wa -k MAC-policy
```

Remediation:

Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor events that modify the system's Mandatory Access Controls.

Example:

```
# printf "  
-w /etc/selinux -p wa -k MAC-policy  
-w /usr/share/selinux -p wa -k MAC-policy  
" >> /etc/audit/rules.d/50-MAC-policy.rules
```

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required.

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot  
required to load rules\n"; fi
```

References:

1. NIST SP 800-53 Rev. 5: AU-3, CM-6

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 <u>Collect Detailed Audit Logs</u> Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.			
v7	5.5 <u>Implement Automated Configuration Monitoring Systems</u> Utilize a Security Content Automation Protocol (SCAP) compliant configuration monitoring system to verify all security configuration elements, catalog approved exceptions, and alert when unauthorized changes occur.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0005	M1022

5.2.3.15 Ensure successful and unsuccessful attempts to use the chcon command are recorded (Automated)

Profile Applicability:

- Level 2 - Server

Description:

The operating system must generate audit records for successful/unsuccessful uses of the `chcon` command.

Rationale:

Without generating audit records that are specific to the security and mission needs of the organization, it would be difficult to establish, correlate, and investigate the events relating to an incident or identify those responsible for one.

Audit records can be generated from various components within the information system (e.g., module or policy filter).

Audit:

64 Bit systems

On disk configuration

Run the following command to check the on disk rules:

```
# {
  UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
  [ -n "${UID_MIN}" ] && awk "/^ *-a *always,exit/ \
&&( / -F *audit!=unset/||/ -F *audit!=-1/||/ -F *audit!=4294967295/) \
&&/ -F *audit>=${UID_MIN}/ \
&&/ -F *perm=x/ \
&&/ -F *path=/usr/bin/chcon/ \
&&( / key= *[-~]* *$/||/ -k *[-~]* *$/)" /etc/audit/rules.d/*.rules \
|| printf "ERROR: Variable 'UID_MIN' is unset.\n"
}
```

Verify the output matches:

```
-a always,exit -F path=/usr/bin/chcon -F perm=x -F audit>=1000 -F audit!=unset
-k perm_chng
```

Running configuration

Run the following command to check loaded rules:

```
# {
  UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
  [ -n "${UID_MIN}" ] && auditctl -l | awk "/^ *-a *always,exit/ \
&&( / -F *audit!=unset/||/ -F *audit!=-1/||/ -F *audit!=4294967295/) \
&&/ -F *audit>=${UID_MIN}/ \
&&/ -F *perm=x/ \
&&/ -F *path=/usr/bin/chcon/ \
&&( / key= *[-~]* *$/||/ -k *[-~]* *$/)" \
|| printf "ERROR: Variable 'UID_MIN' is unset.\n"
}
```

Verify the output matches:

```
-a always,exit -S all -F path=/usr/bin/chcon -F perm=x -F audit>=1000 -F
audit!=-1 -F key=perm_chng
```

32 Bit systems

Follow the same procedures as for 64 bit systems and ignore any entries with b64.

Remediation:

Create audit rules

Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor successful and unsuccessful attempts to use the `chcon` command.

64 Bit systems

Example:

```
# {
  UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
  [ -n "${UID_MIN}" ] && printf "
-a always,exit -F path=/usr/bin/chcon -F perm=x -F auid>=${UID_MIN} -F
auid!=unset -k perm_chng
" >> /etc/audit/rules.d/50-perm_chng.rules || printf "ERROR: Variable
'UID_MIN' is unset.\n"
}
```

Load audit rules

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required.

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

32 Bit systems

Follow the same procedures as for 64 bit systems and ignore any entries with `b64`.

References:

1. NIST SP 800-53 Rev. 5: AU-2, AU-12, SI-5

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0005	M1022

5.2.3.16 Ensure successful and unsuccessful attempts to use the setfacl command are recorded (Automated)

Profile Applicability:

- Level 2 - Server

Description:

The operating system must generate audit records for successful/unsuccessful uses of the `setfacl` command

Rationale:

Without generating audit records that are specific to the security and mission needs of the organization, it would be difficult to establish, correlate, and investigate the events relating to an incident or identify those responsible for one.

Audit records can be generated from various components within the information system (e.g., module or policy filter).

Audit:

64 Bit systems

On disk configuration

Run the following command to check the on disk rules:

```
# {
  UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
  [ -n "${UID_MIN}" ] && awk "/^ *-a *always,exit/ \
    &&( / -F *auid!=unset/||/ -F *auid!=-1/||/ -F *auid!=4294967295/) \
    &&/ -F *auid>=${UID_MIN}/ \
    &&/ -F *perm=x/ \
    &&/ -F *path=/usr/bin/setfacl/ \
    &&( / key= *[-~]* *$/||/ -k *[-~]* *$/)" /etc/audit/rules.d/*.rules \
    || printf "ERROR: Variable 'UID_MIN' is unset.\n"
}
```

Verify the output matches:

```
-a always,exit -F path=/usr/bin/setfacl -F perm=x -F auid>=1000 -F
auid!=unset -k perm_chng
```

Running configuration

Run the following command to check loaded rules:

```
# {
  UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
  [ -n "${UID_MIN}" ] && auditctl -l | awk "/^ *-a *always,exit/ \
    &&( / -F *auid!=unset/||/ -F *auid!=-1/||/ -F *auid!=4294967295/) \
    &&/ -F *auid>=${UID_MIN}/ \
    &&/ -F *perm=x/ \
    &&/ -F *path=/usr/bin/setfacl/ \
    &&( / key= *[-~]* *$/||/ -k *[-~]* *$/)" \
    || printf "ERROR: Variable 'UID_MIN' is unset.\n"
}
```

Verify the output matches:

```
-a always,exit -S all -F path=/usr/bin/setfacl -F perm=x -F auid>=1000 -F
auid!=-1 -F key=perm_chng
```

32 Bit systems

Follow the same procedures as for 64 bit systems and ignore any entries with b64.

Remediation:

Create audit rules

Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor successful and unsuccessful attempts to use the `setfacl` command.

64 Bit systems

Example:

```
# {
  UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
  [ -n "${UID_MIN}" ] && printf "
-a always,exit -F path=/usr/bin/setfacl -F perm=x -F auid>=${UID_MIN} -F
auid!=unset -k perm_chng
" >> /etc/audit/rules.d/50-perm_chng.rules || printf "ERROR: Variable
'UID_MIN' is unset.\n"
}
```

Load audit rules

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required.

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

32 Bit systems

Follow the same procedures as for 64 bit systems and ignore any entries with `b64`.

References:

1. NIST SP 800-53 Rev. 5: AU-2, AU-12, SI-5

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0005	M1022

5.2.3.17 Ensure successful and unsuccessful attempts to use the chacl command are recorded (Automated)

Profile Applicability:

- Level 2 - Server

Description:

The operating system must generate audit records for successful/unsuccessful uses of the `chacl` command

Rationale:

Without generating audit records that are specific to the security and mission needs of the organization, it would be difficult to establish, correlate, and investigate the events relating to an incident or identify those responsible for one.

Audit records can be generated from various components within the information system (e.g., module or policy filter).

Audit:

64 Bit systems

On disk configuration

Run the following command to check the on disk rules:

```
# {
  UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
  [ -n "${UID_MIN}" ] && awk "/^ *-a *always,exit/ \
&&( / -F *audit!=unset/||/ -F *audit!=-1/||/ -F *audit!=4294967295/) \
&&/ -F *audit>=${UID_MIN}/ \
&&/ -F *perm=x/ \
&&/ -F *path=/usr/bin/chacl/ \
&&( / key= *[-~]* *$/||/ -k *[-~]* *$/)" /etc/audit/rules.d/*.rules \
|| printf "ERROR: Variable 'UID_MIN' is unset.\n"
}
```

Verify the output matches:

```
-a always,exit -F path=/usr/bin/chacl -F perm=x -F audit>=1000 -F audit!=unset
-k perm_chng
```

Running configuration

Run the following command to check loaded rules:

```
# {
  UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
  [ -n "${UID_MIN}" ] && auditctl -l | awk "/^ *-a *always,exit/ \
&&( / -F *audit!=unset/||/ -F *audit!=-1/||/ -F *audit!=4294967295/) \
&&/ -F *audit>=${UID_MIN}/ \
&&/ -F *perm=x/ \
&&/ -F *path=/usr/bin/chacl/ \
&&( / key= *[-~]* *$/||/ -k *[-~]* *$/)" \
|| printf "ERROR: Variable 'UID_MIN' is unset.\n"
}
```

Verify the output matches:

```
-a always,exit -S all -F path=/usr/bin/chacl -F perm=x -F audit>=1000 -F
audit!=-1 -F key=perm_chng
```

32 Bit systems

Follow the same procedures as for 64 bit systems and ignore any entries with b64.

Remediation:

Create audit rules

Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor successful and unsuccessful attempts to use the `chac1` command.

64 Bit systems

Example:

```
# {
  UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
  [ -n "${UID_MIN}" ] && printf "
-a always,exit -F path=/usr/bin/chac1 -F perm=x -F auid>=${UID_MIN} -F
auid!=unset -k perm_chng
" >> /etc/audit/rules.d/50-perm_chng.rules || printf "ERROR: Variable
'UID_MIN' is unset.\n"
}
```

Load audit rules

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required.

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

32 Bit systems

Follow the same procedures as for 64 bit systems and ignore any entries with `b64`.

References:

1. NIST SP 800-53 Rev. 5: AU-2, AU-12, SI-5

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0005	M1022

5.2.3.18 Ensure successful and unsuccessful attempts to use the usermod command are recorded (Automated)

Profile Applicability:

- Level 2 - Server

Description:

The operating system must generate audit records for successful/unsuccessful uses of the `usermod` command.

Rationale:

Without generating audit records that are specific to the security and mission needs of the organization, it would be difficult to establish, correlate, and investigate the events relating to an incident or identify those responsible for one.

Audit records can be generated from various components within the information system (e.g., module or policy filter).

Audit:

64 Bit systems

On disk configuration

Run the following command to check the on disk rules:

```
# {
  UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
  [ -n "${UID_MIN}" ] && awk "/^ *-a *always,exit/ \
&&( / -F *auid!=unset/||/ -F *auid!=-1/||/ -F *auid!=4294967295/) \
&&/ -F *auid>=${UID_MIN}/ \
&&/ -F *perm=x/ \
&&/ -F *path=/usr/sbin/usermod/ \
&&( / key= *[-~]* *$/||/ -k *[-~]* *$/)" /etc/audit/rules.d/*.rules \
  || printf "ERROR: Variable 'UID_MIN' is unset.\n"
}
```

Verify the output matches:

```
-a always,exit -F path=/usr/sbin/usermod -F perm=x -F auid>=1000 -F
auid!=unset -k usermod
```

Running configuration

Run the following command to check loaded rules:

```
# {
  UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
  [ -n "${UID_MIN}" ] && auditctl -l | awk "/^ *-a *always,exit/ \
&&( / -F *auid!=unset/||/ -F *auid!=-1/||/ -F *auid!=4294967295/) \
&&/ -F *auid>=${UID_MIN}/ \
&&/ -F *perm=x/ \
&&/ -F *path=/usr/sbin/usermod/ \
&&( / key= *[-~]* *$/||/ -k *[-~]* *$/)" \
  || printf "ERROR: Variable 'UID_MIN' is unset.\n"
}
```

Verify the output matches:

```
-a always,exit -S all -F path=/usr/sbin/usermod -F perm=x -F auid>=1000 -F
auid!=-1 -F key=usermod
```

32 Bit systems

Follow the same procedures as for 64 bit systems and ignore any entries with b64.

Remediation:

Create audit rules

Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor successful and unsuccessful attempts to use the `usermod` command.

64 Bit systems

Example:

```
# {
  UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
  [ -n "${UID_MIN}" ] && printf "
-a always,exit -F path=/usr/sbin/usermod -F perm=x -F auid>=${UID_MIN} -F
auid!=unset -k usermod
" >> /etc/audit/rules.d/50-usermod.rules || printf "ERROR: Variable 'UID_MIN'
is unset.\n"
}
```

Load audit rules

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required.

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

32 Bit systems

Follow the same procedures as for 64 bit systems and ignore any entries with `b64`.

References:

1. NIST SP 800-53 Rev. 5: AU-2, AU-12, SI-5

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 <u>Collect Audit Logs</u> Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0005	M1022

5.2.3.19 Ensure kernel module loading unloading and modification is collected (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Monitor the loading and unloading of kernel modules. All the loading / listing / dependency checking of modules is done by `kmod` via symbolic links.

The following system calls control loading and unloading of modules:

- `init_module` - load a module
- `finit_module` - load a module (used when the overhead of using cryptographically signed modules to determine the authenticity of a module can be avoided)
- `delete_module` - delete a module
- `create_module` - create a loadable module entry
- `query_module` - query the kernel for various bits pertaining to modules

Any execution of the loading and unloading module programs and system calls will trigger an audit record with an identifier of `modules`.

Rationale:

Monitoring the use of all the various ways to manipulate kernel modules could provide system administrators with evidence that an unauthorized change was made to a kernel module, possibly compromising the security of the system.

Audit:

64 Bit systems

On disk configuration

Run the following commands to check the on disk rules:

```
# {
awk '/^ *-a *always,exit/ \
&&/ -F *arch=b[2346]{2}/ \
&&/ -F auid!=unset/||/ -F auid!=-1/||/ -F auid!=4294967295/) \
&&/ -S/ \
&&/init_module/ \
||/finit_module/ \
||/delete_module/ \
||/create_module/ \
||/query_module/) \
&&/ key= *[:-~]* *$/||/ -k *[:-~]* *$/)' /etc/audit/rules.d/*.rules

UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
[ -n "${UID_MIN}" ] && awk "/^ *-a *always,exit/ \
&&/ -F *auid!=unset/||/ -F *auid!=-1/||/ -F *auid!=4294967295/) \
&&/ -F *auid>=${UID_MIN}/ \
&&/ -F *perm=x/ \
&&/ -F *path=/usr/bin/kmod/ \
&&/ key= *[:-~]* *$/||/ -k *[:-~]* *$/)" /etc/audit/rules.d/*.rules \
|| printf "ERROR: Variable 'UID_MIN' is unset.\n"
}
```

Verify the output matches:

```
-a always,exit -F arch=b64 -S
init_module,finit_module,delete_module,create_module,query_module -F
auid>=1000 -F auid!=unset -k kernel_modules
-a always,exit -F path=/usr/bin/kmod -F perm=x -F auid>=1000 -F auid!=unset -
k kernel_modules
```

Running configuration

Run the following command to check loaded rules:


```
# {
auditctl -l | awk '/^ *-a *always,exit/ \
&&/ -F *arch=b[2346]{2}/ \
&&(/ -F auid!=unset/||/ -F auid!=-1/||/ -F auid!=4294967295/) \
&&/ -S/ \
&&(/init_module/ \
||/finit_module/ \
||/delete_module/ \
||/create_module/ \
||/query_module/) \
&&(/ key= *[:-~]* *$/||/ -k *[:-~]* *$/)'

UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
[ -n "${UID_MIN}" ] && auditctl -l | awk "/^ *-a *always,exit/ \
&&(/ -F *auid!=unset/||/ -F *auid!=-1/||/ -F *auid!=4294967295/) \
&&/ -F *auid>=${UID_MIN}/ \
&&/ -F *perm=x/ \
&&/ -F *path=\/usr\/bin\/kmod/ \
&&(/ key= *[:-~]* *$/||/ -k *[:-~]* *$/)" \
|| printf "ERROR: Variable 'UID_MIN' is unset.\n"
}
```

Verify the output includes:

```
-a always,exit -F arch=b64 -S
create_module,init_module,delete_module,query_module,finit_module -F
auid>=1000 -F auid!=-1 -F key=kernel_modules
-a always,exit -S all -F path=/usr/bin/kmod -F perm=x -F auid>=1000 -F
auid!=-1 -F key=kernel_modules
```

Symlink audit

Audit if the symlinks that `kmod` accepts is indeed pointing at it:

```
# S_LINKS=$(ls -l /usr/sbin/lsmmod /usr/sbin/rmmmod /usr/sbin/insmod
/usr/sbin/modinfo /usr/sbin/modprobe /usr/sbin/depmod | grep -v " ->
../bin/kmod" || true) \
&& if [[ "${S_LINKS}" != "" ]]; then printf "Issue with symlinks:
${S_LINKS}\n"; else printf "OK\n"; fi
```

Verify the output states `OK`. If there is a symlink pointing to a different location it should be investigated.

Remediation:

Create audit rules

Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor kernel module modification.

64 Bit systems

Example:

```
# {
UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
[ -n "${UID_MIN}" ] && printf "
-a always,exit -F arch=b64 -S
init_module,fininit_module,delete_module,create_module,query_module -F
audit>=${UID_MIN} -F audit!=unset -k kernel_modules
-a always,exit -F path=/usr/bin/kmod -F perm=x -F audit>=${UID_MIN} -F
audit!=unset -k kernel_modules
" >> /etc/audit/rules.d/50-kernel_modules.rules || printf "ERROR: Variable
'UID_MIN' is unset.\n"
}
```

Load audit rules

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required.

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. NIST SP 800-53 Rev. 5: AU-3, CM-6

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 <u>Collect Detailed Audit Logs</u> Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.		●	●
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.	●	●	●

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.006	TA0004	M1047

5.2.3.20 Ensure the audit configuration is immutable (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Set system audit so that audit rules cannot be modified with `auditctl`. Setting the flag `"-e 2"` forces audit to be put in immutable mode. Audit changes can only be made on system reboot.

Note: This setting will require the system to be rebooted to update the active `auditd` configuration settings.

Rationale:

In immutable mode, unauthorized users cannot execute changes to the audit system to potentially hide malicious activity and then put the audit rules back. Users would most likely notice a system reboot and that could alert administrators of an attempt to make unauthorized audit changes.

Audit:

Run the following command and verify output matches:

```
# grep -Ph -- '^h*-e\h+2\b' /etc/audit/rules.d/*.rules | tail -1
-e 2
```

Remediation:

Edit or create the file `/etc/audit/rules.d/99-finalize.rules` and add the line `-e 2` at the end of the file:

Example:

```
# printf -- "-e 2" >> /etc/audit/rules.d/99-finalize.rules
```

Load audit rules

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required.

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. NIST SP 800-53 Rev. 5: AC-3, AU-3, AU-12, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v8	8.5 <u>Collect Detailed Audit Logs</u> Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.			
v7	6.2 <u>Activate audit logging</u> Ensure that local logging has been enabled on all systems and networking devices.			
v7	6.3 <u>Enable Detailed Logging</u> Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1562, T1562.001	TA0005	

5.2.3.21 *Ensure the running and on disk configuration is the same (Manual)*

Profile Applicability:

- Level 2 - Server

Description:

The Audit system have both on disk and running configuration. It is possible for these configuration settings to differ.

Note: Due to the limitations of `augenrules` and `auditctl`, it is not absolutely guaranteed that loading the rule sets via `augenrules --load` will result in all rules being loaded or even that the user will be informed if there was a problem loading the rules.

Rationale:

Configuration differences between what is currently running and what is on disk could cause unexpected problems or may give a false impression of compliance requirements.

Audit:

Merged rule sets

Ensure that all rules in `/etc/audit/rules.d` have been merged into `/etc/audit/audit.rules`:

```
# augenrules --check
/usr/sbin/augenrules: No change
```

Should there be any drift, run `augenrules --load` to merge and load all rules.

Remediation:

If the rules are not aligned across all three () areas, run the following command to merge and load all rules:

```
# augenrules --load
```

Check if reboot is required.

```
if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then echo "Reboot required
to load rules"; fi
```

References:

1. NIST SP 800-53 Rev. 5: AU-3

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 <u>Collect Detailed Audit Logs</u> Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.			
v7	6.3 <u>Enable Detailed Logging</u> Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.			

5.2.4 Configure auditd file access

Without the capability to restrict which roles and individuals can select which events are audited, unauthorized personnel may be able to prevent the auditing of critical events.

5.2.4.1 Ensure audit log files are mode 0640 or less permissive (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Audit log files contain information about the system and system activity.

Rationale:

Access to audit records can reveal system and configuration data to attackers, potentially compromising its confidentiality.

Audit:

Run the following command to verify audit log files have mode 0640 or less permissive:

```
# [ -f /etc/audit/auditd.conf ] && find "$(dirname $(awk -F "=" '
/^s*log_file/ {print $2}' /etc/audit/auditd.conf | xargs))" -type f \( ! -
perm 600 -a ! -perm 0400 -a ! -perm 0200 -a ! -perm 0000 -a ! -perm 0640 -a !
-perm 0440 -a ! -perm 0040 \) -exec stat -Lc "%n %#a" {} +
```

Nothing should be returned

Remediation:

Run the following command to remove more permissive mode than 0640 from audit log files:

```
# [ -f /etc/audit/auditd.conf ] && find "$(dirname $(awk -F "=" '
/^s*log_file/ {print $2}' /etc/audit/auditd.conf | xargs))" -type f \( ! -
perm 600 -a ! -perm 0400 -a ! -perm 0200 -a ! -perm 0000 -a ! -perm 0640 -a !
-perm 0440 -a ! -perm 0040 \) -exec chmod u-x,g-wx,o-rwx {} +
```

References:

1. NIST SP 800-53 Rev. 5: AU-3

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1083, T1083.000	TA0007	M1022

5.2.4.2 Ensure only authorized users own audit log files (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Audit log files contain information about the system and system activity.

Rationale:

Access to audit records can reveal system and configuration data to attackers, potentially compromising its confidentiality.

Audit:

Run the following command to verify audit log files are owned by the `root` user:

```
# [ -f /etc/audit/auditd.conf ] && find "$(dirname $(awk -F "=" '
/^s*log_file/ {print $2}' /etc/audit/auditd.conf | xargs))" -type f ! -user
root -exec stat -Lc "%n %U" {} +
```

Nothing should be returned

Remediation:

Run the following command to configure the audit log files to be owned by the `root` user:

```
# [ -f /etc/audit/auditd.conf ] && find "$(dirname $(awk -F "=" '
/^s*log_file/ {print $2}' /etc/audit/auditd.conf | xargs))" -type f ! -user
root -exec chown root {} +
```

References:

1. NIST SP 800-53 Rev. 5: AU-3

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1083, T1083.000	TA0007	M1022

5.2.4.3 Ensure only authorized groups are assigned ownership of audit log files (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Audit log files contain information about the system and system activity.

Rationale:

Access to audit records can reveal system and configuration data to attackers, potentially compromising its confidentiality.

Audit:

Run the following command to verify `log_group` parameter is set to either `adm` or `root` in `/etc/audit/auditd.conf`:

```
# grep -Piw -- '^h*log_group\h*=\h*(adm|root)\b' /etc/audit/auditd.conf
```

Verify the output is:

```
log_group = adm
-OR-
log_group = root
```

Using the path of the directory containing the audit logs, determine if the audit log files are owned by the "root" or "adm" group by using the following command:

```
# stat -c "%n %G" "$(dirname $(awk -F"=" ' /^\s*log_file\s*=\s*/ {print $2}' /etc/audit/auditd.conf | xargs))"/* | grep -Pv '^h*\H+\h+(adm|root)\b'
```

Nothing should be returned

Remediation:

Run the following command to configure the audit log files to be owned by `adm` group:

```
# find $(dirname $(awk -F"=" '/^\s*log_file\s*=\s*/ {print $2}'  
/etc/audit/auditd.conf | xargs)) -type f \( ! -group adm -a ! -group root \)  
-exec chgrp adm {} +
```

Run the following command to configure the audit log files to be owned by the `adm` group:

```
# chgrp adm /var/log/audit/
```

Run the following command to set the `log_group` parameter in the audit configuration file to `log_group = adm`:

```
# sed -ri 's/^\s*#\?\s*log_group\s*=\s*\S+(\s*#.*)?.*$/log_group = adm\1/'  
/etc/audit/auditd.conf
```

Run the following command to restart the audit daemon to reload the configuration file:

```
# systemctl restart auditd
```

References:

1. NIST SP 800-53 Rev. 5: AU-3

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1083, T1083.000	TA0007	M1022

5.2.4.4 Ensure the audit log directory is 0750 or more restrictive (Automated)

Profile Applicability:

- Level 2 - Server

Description:

The audit log directory contains audit log files.

Rationale:

Audit information includes all information including: audit records, audit settings and audit reports. This information is needed to successfully audit system activity. This information must be protected from unauthorized modification or deletion. If this information were to be compromised, forensic analysis and discovery of the true source of potentially malicious system activity is impossible to achieve.

Audit:

Run the following command to verify that the audit log directory has a mode of 0750 or less permissive:

```
# stat -Lc "%n %a" "$(dirname $( awk -F"=" ' /^\s*log_file\s*=\s*/ {print $2}' /etc/audit/auditd.conf))" | grep -Pv -- '^h*\H+\h+([0,5,7][0,5]0)'
```

Nothing should be returned

Remediation:

Run the following command to configure the audit log directory to have a mode of "0750" or less permissive:

```
# chmod g-w,o-rwx "$(dirname $( awk -F"=" ' /^\s*log_file\s*=\s*/ {print $2}' /etc/audit/auditd.conf))"
```

Default Value:

750

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1083, T1083.000	TA0007	M1022

5.2.4.5 Ensure audit configuration files are 640 or more restrictive (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Audit configuration files control auditd and what events are audited.

Rationale:

Access to the audit configuration files could allow unauthorized personnel to prevent the auditing of critical events.

Misconfigured audit configuration files may prevent the auditing of critical events or impact the system's performance by overwhelming the audit log. Misconfiguration of the audit configuration files may also make it more difficult to establish and investigate events relating to an incident.

Audit:

Run the following command to verify that the audit configuration files have mode 640 or more restrictive and are owned by the root user and root group:

```
# find /etc/audit/ -type f \( -name '*.conf' -o -name '*.rules' \) -exec stat -Lc "%n %a" {} + | grep -Pv -- '^h*\H+h*([0,2,4,6][0,4]0)\h*$'
```

Nothing should be returned

Remediation:

Run the following command to remove more permissive mode than 0640 from the audit configuration files:

```
# find /etc/audit/ -type f \( -name '*.conf' -o -name '*.rules' \) -exec chmod u-x,g-wx,o-rwx {} +
```

References:

1. NIST SP 800-53 Rev. 5: AU-3

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1083, T1083.000	TA0007	M1022

5.2.4.6 Ensure audit configuration files are owned by root (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Audit configuration files control auditd and what events are audited.

Rationale:

Access to the audit configuration files could allow unauthorized personnel to prevent the auditing of critical events.

Misconfigured audit configuration files may prevent the auditing of critical events or impact the system's performance by overwhelming the audit log. Misconfiguration of the audit configuration files may also make it more difficult to establish and investigate events relating to an incident.

Audit:

Run the following command to verify that the audit configuration files have mode 640 or more restrictive and are owned by the root user and root group:

```
# find /etc/audit/ -type f \( -name '*.conf' -o -name '*.rules' \) ! -user root
```

Nothing should be returned

Remediation:

Run the following command to change ownership to `root` user:

```
# find /etc/audit/ -type f \( -name '*.conf' -o -name '*.rules' \) ! -user root -exec chown root {} +
```

References:

1. NIST SP 800-53 Rev. 5: AU-3

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1083, T1083.000	TA0007	M1022

5.2.4.7 Ensure audit configuration files belong to group root (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Audit configuration files control auditd and what events are audited.

Rationale:

Access to the audit configuration files could allow unauthorized personnel to prevent the auditing of critical events.

Misconfigured audit configuration files may prevent the auditing of critical events or impact the system's performance by overwhelming the audit log. Misconfiguration of the audit configuration files may also make it more difficult to establish and investigate events relating to an incident.

Audit:

Run the following command to verify that the audit configuration files have mode 640 or more restrictive and are owned by the root user and root group:

```
# find /etc/audit/ -type f \( -name '*.conf' -o -name '*.rules' \) ! -group root
```

Nothing should be returned

Remediation:

Run the following command to change group to `root`:

```
# find /etc/audit/ -type f \( -name '*.conf' -o -name '*.rules' \) ! -group root -exec chgrp root {} +
```

References:

1. NIST SP 800-53 Rev. 5: AU-3

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1083, T1083.000	TA0007	M1022

5.2.4.8 Ensure audit tools are 755 or more restrictive (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Audit tools include, but are not limited to, vendor-provided and open source audit tools needed to successfully view and manipulate audit information system activity and records. Audit tools include custom queries and report generators.

Rationale:

Protecting audit information includes identifying and protecting the tools used to view and manipulate log data. Protecting audit tools is necessary to prevent unauthorized operation on audit information.

Audit:

Run the following command to verify the audit tools have mode 755 or more restrictive, are owned by the root user and group root:

```
# stat -c "%n %a" /sbin/auditctl /sbin/aureport /sbin/ausearch /sbin/autrace  
/sbin/auditd /sbin/augenrules | grep -Pv -- '^h*\H+\h+([0-  
7][0,1,4,5][0,1,4,5])\h*$'
```

Nothing should be returned

Remediation:

Run the following command to remove more permissive mode from the audit tools:

```
# chmod go-w /sbin/auditctl /sbin/aureport /sbin/ausearch /sbin/autrace  
/sbin/auditd /sbin/augenrules
```

References:

1. NIST SP 800-53 Rev. 5: AU-3

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1083, T1083.000	TA0007	M1022

5.2.4.9 Ensure audit tools are owned by root (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Audit tools include, but are not limited to, vendor-provided and open source audit tools needed to successfully view and manipulate audit information system activity and records. Audit tools include custom queries and report generators.

Rationale:

Protecting audit information includes identifying and protecting the tools used to view and manipulate log data. Protecting audit tools is necessary to prevent unauthorized operation on audit information.

Audit:

Run the following command to verify the audit tools have mode 755 or more restrictive, are owned by the root user and group root:

```
# stat -c "%n %U" /sbin/auditctl /sbin/aureport /sbin/ausearch /sbin/autrace  
/sbin/auditd /sbin/augenrules | grep -Pv -- '^h*\H+h+root\h*$'
```

Nothing should be returned

Remediation:

Run the following command to change the owner of the audit tools to the root user:

```
# chown root /sbin/auditctl /sbin/aureport /sbin/ausearch /sbin/autrace  
/sbin/auditd /sbin/augenrules
```

References:

1. NIST SP 800-53 Rev. 5: AU-3

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1083, T1083.000	TA0007	M1022

5.2.4.10 Ensure audit tools belong to group root (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Audit tools include, but are not limited to, vendor-provided and open source audit tools needed to successfully view and manipulate audit information system activity and records. Audit tools include custom queries and report generators.

Rationale:

Protecting audit information includes identifying and protecting the tools used to view and manipulate log data. Protecting audit tools is necessary to prevent unauthorized operation on audit information.

Audit:

Run the following command to verify the audit tools have mode 755 or more restrictive, are owned by the root user and group root:

```
# stat -c "%n %a %U %G" /sbin/auditctl /sbin/aureport /sbin/ausearch  
/sbin/autrace /sbin/auditd /sbin/augenrules | grep -Pv -- '^h*\H+h+([0-  
7][0,1,4,5][0,1,4,5])\h+root\h+root\h*$'
```

Nothing should be returned

Remediation:

Run the following command to remove more permissive mode from the audit tools:

```
# chmod go-w /sbin/auditctl /sbin/aureport /sbin/ausearch /sbin/autrace  
/sbin/auditd /sbin/augenrules
```

Run the following command to change owner and group of the audit tools to `root` user and group:

```
# chown root:root /sbin/auditctl /sbin/aureport /sbin/ausearch /sbin/autrace  
/sbin/auditd /sbin/augenrules
```

References:

1. NIST SP 800-53 Rev. 5: AU-3

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002, T1083, T1083.000	TA0007	M1022

5.3 Ensure logrotate is configured (Manual)

Profile Applicability:

- Level 1 - Server

Description:

The system includes the capability of rotating log files regularly to avoid filling up the system with logs or making the logs unmanageably large. The file `/etc/logrotate.d/syslog` is the configuration file used to rotate log files created by `syslog` or `rsyslog`.

Rationale:

By keeping the log files smaller and more manageable, a system administrator can easily archive these files to another system and spend less time looking through inordinately large log files.

Audit:

Review `/etc/logrotate.conf` and `/etc/logrotate.d/*` and verify logs are rotated according to site policy.

Remediation:

Edit `/etc/logrotate.conf` and `/etc/logrotate.d/*` to ensure logs are rotated according to site policy.

References:

1. NIST SP 800-53 Rev. 5: AU-8

Additional Information:

If no `maxage` setting is set for `logrotate` a situation can occur where `logrotate` is interrupted and fails to delete rotated log files. It is recommended to set this to a value greater than the longest any log file should exist on your system to ensure that any such log file is removed but standard rotation settings are not overridden.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.3 <u>Ensure Adequate Audit Log Storage</u> Ensure that logging destinations maintain adequate storage to comply with the enterprise's audit log management process.			
v7	6.4 <u>Ensure adequate storage for logs</u> Ensure that all systems that store logs have adequate storage space for the logs generated.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1070, T1070.002	TA0040	M1022

6 System Maintenance

Recommendations in this section are intended as maintenance and are intended to be checked on a frequent basis to ensure system stability. Many recommendations do not have quick remediations and require investigation into the cause and best fix available and may indicate an attempted breach of system security.

6.1 System File Permissions

This section provides guidance on securing aspects of system files and directories.

6.1.1 Ensure permissions on /etc/passwd are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `/etc/passwd` file contains user account information that is used by many system utilities and therefore must be readable for these utilities to operate.

Rationale:

It is critical to ensure that the `/etc/passwd` file is protected from unauthorized write access. Although it is protected by default, the file permissions could be changed either inadvertently or through malicious actions.

Audit:

Run the following command to verify `/etc/passwd` is mode 644 or more restrictive, Uid is 0/root and Gid is 0/root:

```
# stat -Lc "%n %a %u/%U %g/%G" /etc/passwd
/etc/passwd 644 0/root 0/root
```

Remediation:

Run the following commands to remove excess permissions, set owner, and set group on `/etc/passwd`:

```
# chmod u-x,go-wx /etc/passwd
# chown root:root /etc/passwd
```

Default Value:

`/etc/passwd 644 0/root 0/root`

References:

1. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	16.4 <u>Encrypt or Hash all Authentication Credentials</u> Encrypt or hash with a salt all authentication credentials when stored.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1003, T1003.008, T1222, T1222.002	TA0005	M1022

6.1.2 Ensure permissions on /etc/passwd are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `/etc/passwd` file contains user account information that is used by many system utilities and therefore must be readable for these utilities to operate.

Rationale:

It is critical to ensure that the `/etc/passwd` file is protected from unauthorized write access. Although it is protected by default, the file permissions could be changed either inadvertently or through malicious actions.

Audit:

Run the following command to verify `/etc/passwd` is mode 644 or more restrictive, Uid is 0/root and Gid is 0/root:

```
# stat -Lc "%n %a %u/%U %g/%G" /etc/passwd
/etc/passwd 644 0/root 0/root
```

Remediation:

Run the following commands to remove excess permissions, set owner, and set group on `/etc/passwd`:

```
# chmod u-x,go-wx /etc/passwd
# chown root:root /etc/passwd
```

Default Value:

`/etc/passwd 644 0/root 0/root`

References:

1. NIST SP 800-53 Rev. 5: AC-3. MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	16.4 <u>Encrypt or Hash all Authentication Credentials</u> Encrypt or hash with a salt all authentication credentials when stored.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1003, T1003.008, T1222, T1222.002	TA0005	M1022

6.1.3 Ensure permissions on /etc/passwd- are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The /etc/passwd- file contains backup user account information.

Rationale:

It is critical to ensure that the /etc/passwd- file is protected from unauthorized access. Although it is protected by default, the file permissions could be changed either inadvertently or through malicious actions.

Audit:

Run the following command to verify /etc/passwd- is mode 644 or more restrictive, Uid is 0/root and Gid is 0/root:

```
# stat -Lc "%n %a %u/%U %g/%G" /etc/passwd-  
  
/etc/passwd- 644 0/root 0/root
```

Remediation:

Run the following commands to remove excess permissions, set owner, and set group on /etc/passwd-:

```
# chmod u-x,go-wx /etc/passwd-  
# chown root:root /etc/passwd-
```

Default Value:

/etc/passwd- 644 0/root 0/root

References:

1. NIST SP 800-53 Rev. 5: AC-3. MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	16.4 <u>Encrypt or Hash all Authentication Credentials</u> Encrypt or hash with a salt all authentication credentials when stored.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1003, T1003.008, T1222, T1222.002	TA0005	M1022

6.1.4 Ensure permissions on /etc/group are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `/etc/group` file contains a list of all the valid groups defined in the system. The command below allows read/write access for root and read access for everyone else.

Rationale:

The `/etc/group` file needs to be protected from unauthorized changes by non-privileged users, but needs to be readable as this information is used with many non-privileged programs.

Audit:

Run the following command to verify `/etc/group` is mode 644 or more restrictive, Uid is 0/root and Gid is 0/root:

```
# stat -Lc "%n %a %u/%U %g/%G" /etc/group
/etc/group 644 0/root 0/root
```

Remediation:

Run the following commands to remove excess permissions, set owner, and set group on `/etc/group`:

```
# chmod u-x,go-wx /etc/group
# chown root:root /etc/group
```

Default Value:

`/etc/group 644 0/root 0/root`

References:

1. NIST SP 800-53 Rev. 5: AC-3. MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	16.4 <u>Encrypt or Hash all Authentication Credentials</u> Encrypt or hash with a salt all authentication credentials when stored.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1003, T1003.008, T1222, T1222.002	TA0005	M1022

6.1.5 Ensure permissions on /etc/group- are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The /etc/group- file contains a backup list of all the valid groups defined in the system.

Rationale:

It is critical to ensure that the /etc/group- file is protected from unauthorized access. Although it is protected by default, the file permissions could be changed either inadvertently or through malicious actions.

Audit:

Run the following command to verify /etc/group- is mode 644 or more restrictive, Uid is 0/root and Gid is 0/root:

```
# stat -Lc "%n %a %u/%U %g/%G" /etc/group-  
/etc/group- 644 0/root 0/root
```

Remediation:

Run the following commands to remove excess permissions, set owner, and set group on /etc/group-:

```
# chmod u-x,go-wx /etc/group-  
# chown root:root /etc/group-
```

Default Value:

/etc/group- 644 0/root 0/root

References:

1. NIST SP 800-53 Rev. 5: AC-3. MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	16.4 <u>Encrypt or Hash all Authentication Credentials</u> Encrypt or hash with a salt all authentication credentials when stored.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1003, T1003.008, T1222, T1222.002	TA0005	M1022

6.1.6 Ensure permissions on /etc/shadow are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `/etc/shadow` file is used to store the information about user accounts that is critical to the security of those accounts, such as the hashed password and other security information.

Rationale:

If attackers can gain read access to the `/etc/shadow` file, they can easily run a password cracking program against the hashed password to break it. Other security information that is stored in the `/etc/shadow` file (such as expiration) could also be useful to subvert the user accounts.

Audit:

Run the following command to verify `/etc/shadow` is mode 000, Uid is 0/root and Gid is 0/root:

```
# stat -Lc "%n %a %u/%U %g/%G" /etc/shadow
/etc/shadow 0 0/root 0/root
```

Remediation:

Run the following commands to set mode, owner, and group on `/etc/shadow`:

```
# chown root:root /etc/shadow
# chmod 0000 /etc/shadow
```

Default Value:

`/etc/shadow 0 0/root 0/root`

References:

1. NIST SP 800-53 Rev. 5: AC-3. MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	16.4 <u>Encrypt or Hash all Authentication Credentials</u> Encrypt or hash with a salt all authentication credentials when stored.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1003, T1003.008, T1222, T1222.002	TA0005	M1022

6.1.7 Ensure permissions on /etc/shadow- are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `/etc/shadow-` file is used to store backup information about user accounts that is critical to the security of those accounts, such as the hashed password and other security information.

Rationale:

It is critical to ensure that the `/etc/shadow-` file is protected from unauthorized access. Although it is protected by default, the file permissions could be changed either inadvertently or through malicious actions.

Audit:

Run the following command to verify `/etc/shadow-` is mode 000, Uid is 0/root and Gid is 0/root:

```
# stat -Lc "%n %a %u/%U %g/%G" /etc/shadow-  
/etc/shadow- 0 0/root 0/root
```

Remediation:

Run the following commands to set mode, owner, and group on `/etc/shadow-`:

```
# chown root:root /etc/shadow-  
# chmod 0000 /etc/shadow-
```

Default Value:

`/etc/shadow- 0 0/root 0/root`

References:

1. NIST SP 800-53 Rev. 5: AC-3. MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	16.4 <u>Encrypt or Hash all Authentication Credentials</u> Encrypt or hash with a salt all authentication credentials when stored.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1003, T1003.008, T1222, T1222.002	TA0005	M1022

6.1.8 Ensure permissions on /etc/gshadow are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `/etc/gshadow` file is used to store the information about groups that is critical to the security of those accounts, such as the hashed password and other security information.

Rationale:

If attackers can gain read access to the `/etc/gshadow` file, they can easily run a password cracking program against the hashed password to break it. Other security information that is stored in the `/etc/gshadow` file (such as group administrators) could also be useful to subvert the group.

Audit:

Run the following command to verify `/etc/gshadow` is mode `000`, Uid is `0/root` and Gid is `0/root`:

```
# stat -Lc "%n %a %u/%U %g/%G" /etc/gshadow
/etc/gshadow 0 0/root 0/root
```

Remediation:

Run the following commands to set mode, owner, and group on `/etc/gshadow`:

```
# chown root:root /etc/gshadow
# chmod 0000 /etc/gshadow
```

Default Value:

`/etc/gshadow 0 0/root 0/root`

References:

1. NIST SP 800-53 Rev. 5: AC-3. MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	16.4 <u>Encrypt or Hash all Authentication Credentials</u> Encrypt or hash with a salt all authentication credentials when stored.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1003, T1003.008, T1222, T1222.002	TA0005	M1022

6.1.9 Ensure permissions on /etc/gshadow- are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `/etc/gshadow-` file is used to store backup information about groups that is critical to the security of those accounts, such as the hashed password and other security information.

Rationale:

It is critical to ensure that the `/etc/gshadow-` file is protected from unauthorized access. Although it is protected by default, the file permissions could be changed either inadvertently or through malicious actions.

Audit:

Run the following command to verify `/etc/gshadow-` is mode 000, Uid is 0/root and Gid is 0/root:

```
# stat -Lc "%n %a %u/%U %g/%G" /etc/gshadow-  
/etc/gshadow- 0 0/root 0/root
```

Remediation:

Run the following commands to set mode, owner, and group on `/etc/gshadow-`:

```
# chown root:root /etc/gshadow-  
# chmod 0000 /etc/gshadow-
```

Default Value:

`/etc/gshadow- 0 0/root 0/root`

References:

1. NIST SP 800-53 Rev. 5: AC-3. MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	16.4 <u>Encrypt or Hash all Authentication Credentials</u> Encrypt or hash with a salt all authentication credentials when stored.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1003, T1003.008, T1222, T1222.002	TA0005	M1022

6.1.10 Audit system file permissions (Manual)

Profile Applicability:

- Level 2 - Server

Description:

The RPM package manager has a number of useful options. One of these, the `-v` for RPM option, can be used to verify that system packages are correctly installed. The `-v` option can be used to verify a particular package or to verify all system packages. If no output is returned, the package is installed correctly. The following table describes the meaning of output from the verify option:

Code	Meaning
S	File size differs.
M	File mode differs (includes permissions and file type).
5	The MD5 checksum differs.
D	The major and minor version numbers differ on a device file.
L	A mismatch occurs in a link.
U	The file ownership differs.
G	The file group owner differs.
T	The file time (mtime) differs.

The `rpm -qf` command can be used to determine which package a particular file belongs to. For example, the following commands determines which package the `/bin/bash` file belongs to:

```
# rpm -qf /bin/bash
bash-4.1.2-29.el6.x86_64
# rpm -S /bin/bash
bash: /bin/bash
```

To verify the settings for the package that controls the `/bin/bash` file, run the following:

```
# rpm -V bash-4.1.2-29.el6.x86_64
.M..... /bin/bash
# rpm --verify bash
??5????? c /etc/bash.bashrc
```

Note that you can feed the output of the `rpm -qf` command to the `rpm -v` command:

```
# rpm -V `rpm -qf /etc/passwd`
.M..... c /etc/passwd
S.5....T c /etc/printcap
```

Rationale:

It is important to confirm that packaged system files and directories are maintained with the permissions they were intended to have from the OS vendor.

Audit:

Run the following command to review all installed packages. Note that this may be very time consuming and may be best scheduled via the `cron` utility. It is recommended that the output of this command be redirected to a file that can be reviewed later.

```
# rpm -Va --nomtime --nosize --nomd5 --nolinkto > <filename>
```

Remediation:

Correct any discrepancies found and rerun the audit until output is clean or risk is mitigated or accepted.

References:

1. http://docs.fedoraproject.org/en-US/Fedora_Draft_Documentation/0.1/html/RPM_Guide/index.html
2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5, AC-3, MP-2

Additional Information:

Since packages and important files may change with new updates and releases, it is recommended to verify everything, not just a finite list of files. This can be a time consuming task and results may depend on site policy therefore it is not a scorable benchmark item, but is provided for those interested in additional security measures.

Some of the recommendations of this benchmark alter the state of files audited by this recommendation. The audit command will alert for all changes to a file's permissions even if the new state is more secure than the default.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1222	TA0005	M1022

6.1.11 Ensure world writable files and directories are secured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

World writable files are the least secure. Data in world-writable files can be modified and compromised by any user on the system. World writable files may also indicate an incorrectly written script or program that could potentially be the cause of a larger compromise to the system's integrity. See the `chmod(2)` man page for more information.

Setting the sticky bit on world writable directories prevents users from deleting or renaming files in that directory that are not owned by them.

Rationale:

Data in world-writable files can be modified and compromised by any user on the system. World writable files may also indicate an incorrectly written script or program that could potentially be the cause of a larger compromise to the system's integrity.

This feature prevents the ability to delete or rename files in world writable directories (such as `/tmp`) that are owned by another user.

Audit:

Run the following script to verify:

- No world writable files exist
- No world writable directories without the sticky bit exist

```
#!/usr/bin/env bash

{
    l_output="" l_output2=""
    l_smask='01000'
    a_path=(); a_arr=(); a_file=(); a_dir=() # Initialize arrays
    a_path=(! -path "/run/user/*" -a ! -path "/proc/*" -a ! -path
    "*/containerd/*" -a ! -path "*/kubelet/pods/*" -a ! -path
    "/sys/kernel/security/apparmor/*" -a ! -path "/snap/*" -a ! -path
    "/sys/fs/cgroup/memory/*")
    while read -r l_bfs; do
        a_path+=( -a ! -path "$l_bfs/*")
        done < <(findmnt -Dkerno fstype,target | awk ' $1 ~ /^s*(nfs|proc|smb)/
{print $2}')
        # Populate array with files that will possibly fail one of the audits
        while IFS= read -r -d $'\0' l_file; do
            [ -e "$l_file" ] && a_arr+=("$l_file")
        done < <(find / \( "${a_path[@]}" \) \( -type f -o -type d \) -perm -0002
-print0 2>/dev/null)
        while IFS="" read -r l_fname l_mode; do # Test files in the array
            [ -f "$l_fname" ] && a_file+=("$l_fname") # Add WR files
            if [ -d "$l_fname" ]; then # Add directories w/o sticky bit
                [ ! $( ($l_mode & $l_smask) -gt 0 ) ] && a_dir+=("$l_fname")
            fi
            done < <(printf '%s\n' "${a_arr[@]}")
            if ! (( ${#a_file[@]} > 0 )); then
                l_output="$l_output\n - No world writable files exist on the local
filesystem."
            else
                l_output2="$l_output2\n - There are \"$(printf '%s' "${a_file[@]}")\"
World writable files on the system.\n - The following is a list of World
writable files:\n$(printf '%s\n' "${a_file[@]}")\n - end of list\n"
            fi
            if ! (( ${#a_dir[@]} > 0 )); then
                l_output="$l_output\n - Sticky bit is set on world writable
directories on the local filesystem."
            else
                l_output2="$l_output2\n - There are \"$(printf '%s' "${a_dir[@]}")\"
World writable directories without the sticky bit on the system.\n - The
following is a list of World writable directories without the sticky
bit:\n$(printf '%s\n' "${a_dir[@]}")\n - end of list\n"
            fi
            unset a_path; unset a_arr; unset a_file; unset a_dir # Remove arrays
            # If l_output2 is empty, we pass
            if [ -z "$l_output2" ]; then
                echo -e "\n- Audit Result:\n ** PASS **\n - * Correctly configured *
:\n$l_output\n"
            else
                echo -e "\n- Audit Result:\n ** FAIL **\n - * Reasons for audit
failure * :\n$l_output2"
                [ -n "$l_output" ] && echo -e "- * Correctly configured *
:\n$l_output\n"
            fi
        }
}
```

Note: On systems with a large number of files and/or directories, this audit may be a long running process

Remediation:

- World Writable Files:
 - It is recommended that write access is removed from `other` with the command (`chmod o-w <filename>`), but always consult relevant vendor documentation to avoid breaking any application dependencies on a given file.
- World Writable Directories:
 - Set the sticky bit on all world writable directories with the command (`chmod a+t <directory_name>`)

Run the following script to:

- Remove other write permission from any world writable files
- Add the sticky bit to all world writable directories


```
#!/usr/bin/env bash

{
    l_smask='01000'
    a_path=(); a_arr=() # Initialize array
    a_path=(! -path "/run/user/*" -a ! -path "/proc/*" -a ! -path
"/sys/kernel/security/apparmor/*" -a ! -path "/snap/*" -a ! -path
"/sys/fs/cgroup/memory/*")
    while read -r l_bfs; do
        a_path+=( -a ! -path "$l_bfs/*")
        done < <(findmnt -Dkerno fstype,target | awk '$1 ~ /^s*(nfs|proc|smb)/
{print $2}')
        # Populate array with files
        while IFS= read -r -d $'\0' l_file; do
            [ -e "$l_file" ] && a_arr+=("$l_file")
        done < <(find / \( "${a_path[@]}" \) \( -type f -o -type d \) -perm -0002
-print0 2>/dev/null)
        while IFS="^" read -r l_fname l_mode; do # Test files in the array
            if [ -f "$l_fname" ]; then # Remove excess permissions from WW files
                echo -e " - File: \"$l_fname\" is mode: \"$l_mode\" \n - removing
write permission on \"$l_fname\" from \"other\""
                chmod o-w "$l_fname"
            fi
            if [ -d "$l_fname" ]; then
                if [ ! $(( $l_mode & $l_smask )) -gt 0 ]; then # Add sticky bit
                    echo -e " - Directory: \"$l_fname\" is mode: \"$l_mode\" and
doesn't have the sticky bit set \n - Adding the sticky bit"
                    chmod a+t "$l_fname"
                fi
            fi
        done < <(printf '%s\n' "${a_arr[@]}")
        unset a_path; unset a_arr # Remove array
    }
}
```

References:

1. NIST SP 800-53 Rev. 5: AC-3. MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1222, T1222.002, T1548	TA0004, TA0005	M1022, M1028

6.1.12 Ensure no unowned or ungrouped files or directories exist (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Administrators may delete users or groups from the system and neglect to remove all files and/or directories owned by those users or groups.

Rationale:

A new user or group who is assigned a deleted user's user ID or group ID may then end up "owning" a deleted user or group's files, and thus have more access on the system than was intended.

Audit:

Run the following script to verify no unowned or ungrouped files or directories exist:

```
#!/usr/bin/env bash

{
    l_output="" l_output2=""
    a_path=(); a_arr=(); a_nouser=(); a_nogroup=() # Initialize arrays
    a_path=(! -path "/run/user/*" -a ! -path "/proc/*" -a ! -path
    "*/containerd/*" -a ! -path "*/kubelet/pods/*")
    while read -r l_bfs; do
        a_path+=( -a ! -path "$l_bfs/*")
        done < <(findmnt -Dkerno fstype,target | awk ' $1 ~ /^\/s*(nfs|proc|smb)/
    {print $2}')
        while IFS= read -r -d $'\0' l_file; do
            [ -e "$l_file" ] && a_arr+=("$ (stat -Lc '%n^U^G' "$l_file")") && echo
    "Adding: $l_file"
            done < <(find / \( "${a_path[@]}" \) \( -type f -o -type d \) \( -nouser -
    o -nogroup \) -print0 2> /dev/null)
            while IFS="^" read -r l_fname l_user l_group; do # Test files in the array
                [ "$l_user" = "UNKNOWN" ] && a_nouser+=("$l_fname")
                [ "$l_group" = "UNKNOWN" ] && a_nogroup+=("$l_fname")
            done <<< "$ (printf '%s\n' "${a_arr[@]}")"
            if ! (( ${#a_nouser[@]} > 0 )); then
                l_output="$l_output\n - No unowned files or directories exist on the
    local filesystem."
            else
                l_output2="$l_output2\n - There are \"$(printf '%s'
    "${#a_nouser[@]})\" unowned files or directories on the system.\n - The
    following is a list of unowned files and/or directories:\n$(printf '%s\n'
    "${a_nouser[@]})\" \n - end of list"
            fi
            if ! (( ${#a_nogroup[@]} > 0 )); then
                l_output="$l_output\n - No ungrouped files or directories exist on the
    local filesystem."
            else
                l_output2="$l_output2\n - There are \"$(printf '%s'
    "${#a_nogroup[@]})\" ungrouped files or directories on the system.\n - The
    following is a list of ungrouped files and/or directories:\n$(printf '%s\n'
    "${a_nogroup[@]})\" \n - end of list"
            fi
            unset a_path; unset a_arr ; unset a_nouser; unset a_nogroup # Remove
    arrays
            if [ -z "$l_output2" ]; then # If l_output2 is empty, we pass
                echo -e "\n- Audit Result:\n ** PASS **\n - * Correctly configured *
    :\n$l_output\n"
            else
                echo -e "\n- Audit Result:\n ** FAIL **\n - * Reasons for audit
    failure * :\n$l_output2"
                [ -n "$l_output" ] && echo -e "\n- * Correctly configured *
    :\n$l_output\n"
            fi
        }
    }
```

Note: On systems with a large number of files and/or directories, this audit may be a long running process

Remediation:

Remove or set ownership and group ownership of these files and/or directories to an active user on the system as appropriate.

References:

1. NIST SP 800-53 Rev. 5: AC-3. MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v8	14.6 <u>Train Workforce Members on Recognizing and Reporting Security Incidents</u> Train workforce members to be able to recognize a potential incident and be able to report such an incident.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1222, T1222.002	TA0007	M1022

6.1.13 Ensure SUID and SGID files are reviewed (Manual)

Profile Applicability:

- Level 1 - Server

Description:

The owner of a file can set the file's permissions to run with the owner's or group's permissions, even if the user running the program is not the owner or a member of the group. The most common reason for a SUID or SGID program is to enable users to perform functions (such as changing their password) that require root privileges.

Rationale:

There are valid reasons for SUID and SGID programs, but it is important to identify and review such programs to ensure they are legitimate. Review the files returned by the action in the audit section and check to see if system binaries have a different checksum than what from the package. This is an indication that the binary may have been replaced.

Audit:

Run the following script to generate a list of SUID and SGID files:

```
#!/usr/bin/env bash

{
    l_output="" l_output2=""
    a_arr=(); a_suid=(); a_sgid=() # initialize arrays
    # Populate array with files that will possibly fail one of the audits
    while read -r l_mpname; do
        while IFS= read -r -d $'\0' l_file; do
            [ -e "$l_file" ] && a_arr+=("$$(stat -Lc '%n^%#a' "$l_file")")
            done << $(find "$l_mpname" -xdev -not -path "/run/user/*" -type f \( -
perm -2000 -o -perm -4000 \) -print0)
            done <<< "$(findmnt -Derno target)"
            # Test files in the array
            while IFS="^" read -r l_fname l_mode; do
                if [ -f "$l_fname" ]; then
                    l_suid_mask="04000"; l_sgid_mask="02000"
                    [ $(( $l_mode & $l_suid_mask )) -gt 0 ] && a_suid+=("$l_fname")
                    [ $(( $l_mode & $l_sgid_mask )) -gt 0 ] && a_sgid+=("$l_fname")
                fi
            done <<< "$(printf '%s\n' "${a_arr[@]}")"
            if ! (( ${#a_suid[@]} > 0 )); then
                l_output="$l_output\n - There are no SUID files exist on the system"
            else
                l_output2="$l_output2\n - List of \"$(printf '%s' "${a_suid[@]}")\"
SUID executable files:\n$(printf '%s\n' "${a_suid[@]}")\n - end of list -\n"
            fi
            if ! (( ${#a_sgid[@]} > 0 )); then
                l_output="$l_output\n - There are no SGID files exist on the system"
            else
                l_output2="$l_output2\n - List of \"$(printf '%s' "${a_sgid[@]}")\"
SGID executable files:\n$(printf '%s\n' "${a_sgid[@]}")\n - end of list -\n"
            fi
            [ -n "$l_output2" ] && l_output2="$l_output2\n- Review the preceding
list(s) of SUID and/or SGID files to\n- ensure that no rogue programs have
been introduced onto the system.\n"
            unset a_arr; unset a_suid; unset a_sgid # Remove arrays
            # If l_output2 is empty, Nothing to report
            if [ -z "$l_output2" ]; then
                echo -e "\n- Audit Result:\n$l_output\n"
            else
                echo -e "\n- Audit Result:\n$l_output2\n"
                [ -n "$l_output" ] && echo -e "$l_output\n"
            fi
        fi
    done
}
```

Note: on systems with a large number of files, this may be a long running process

Remediation:

Ensure that no rogue SUID or SGID programs have been introduced into the system. Review the files returned by the action in the Audit section and confirm the integrity of these binaries.

References:

1. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5, AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1548, T1548.001	TA0004	M1028

6.2 Local User and Group Settings

This section provides guidance on securing aspects of the local users and groups.

Note: The recommendations in this section check local users and groups. Any users or groups from other sources such as LDAP will not be audited. In a domain environment, similar checks should be performed against domain users and groups.

6.2.1 Ensure accounts in `/etc/passwd` use shadowed passwords (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Local accounts can use shadowed passwords. With shadowed passwords, the passwords are saved in shadow password file, `/etc/shadow`, encrypted by a salted one-way hash. Accounts with a shadowed password have an `x` in the second field in `/etc/passwd`.

Rationale:

The `/etc/passwd` file also contains information like user ID's and group ID's that are used by many system programs. Therefore, the `/etc/passwd` file must remain world readable. In spite of encoding the password with a randomly-generated one-way hash function, an attacker could still break the system if they got access to the `/etc/passwd` file. This can be mitigated by using shadowed passwords, thus moving the passwords in the `/etc/passwd` file to `/etc/shadow`. The `/etc/shadow` file is set so only root will be able to read and write. This helps mitigate the risk of an attacker gaining access to the encoded passwords with which to perform a dictionary attack.

Note:

- All accounts must have passwords or be locked to prevent the account from being used by an unauthorized user.
- A user account with an empty second field in `/etc/passwd` allows the account to be logged into by providing only the username.

Audit:

Run the following command and verify that no output is returned:

```
# awk -F: '($2 != "x" ) { print $1 " is not set to shadowed passwords " }' /etc/passwd
```

Remediation:

Run the following command to set accounts to use shadowed passwords:

```
# sed -e 's/^\([a-zA-Z0-9_]*\) :[^:]*:/\1:x:/' -i /etc/passwd
```

Investigate to determine if the account is logged in and what it is being used for, to determine if it needs to be forced off.

References:

1. NIST SP 800-53 Rev. 5: IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.11 <u>Encrypt Sensitive Data at Rest</u> Encrypt sensitive data at rest on servers, applications, and databases containing sensitive data. Storage-layer encryption, also known as server-side encryption, meets the minimum requirement of this Safeguard. Additional encryption methods may include application-layer encryption, also known as client-side encryption, where access to the data storage device(s) does not permit access to the plain-text data.			
v7	16.4 <u>Encrypt or Hash all Authentication Credentials</u> Encrypt or hash with a salt all authentication credentials when stored.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1003, T1003.008	TA0003	M1027

6.2.2 Ensure /etc/shadow password fields are not empty (Automated)

Profile Applicability:

- Level 1 - Server

Description:

An account with an empty password field means that anybody may log in as that user without providing a password.

Rationale:

All accounts must have passwords or be locked to prevent the account from being used by an unauthorized user.

Audit:

Run the following command and verify that no output is returned:

```
# awk -F: '($2 == "" ) { print $1 " does not have a password "}' /etc/shadow
```

Remediation:

If any accounts in the `/etc/shadow` file do not have a password, run the following command to lock the account until it can be determined why it does not have a password:

```
# passwd -l <username>
```

Also, check to see if the account is logged in and investigate what it is being used for to determine if it needs to be forced off.

References:

1. NIST SP 800-53 Rev. 5: IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.2 Use Unique Passwords Use unique passwords for all enterprise assets. Best practice implementation includes, at a minimum, an 8-character password for accounts using MFA and a 14-character password for accounts not using MFA.			
v7	4.4 Use Unique Passwords Where multi-factor authentication is not supported (such as local administrator, root, or service accounts), accounts will use passwords that are unique to that system.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1078, T1078.001, T1078.003	TA0003	M1027

6.2.3 Ensure all groups in /etc/passwd exist in /etc/group (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Over time, system administration errors and changes can lead to groups being defined in /etc/passwd but not in /etc/group .

Rationale:

Groups defined in the /etc/passwd file but not in the /etc/group file pose a threat to system security since group permissions are not properly managed.

Audit:

Run the following script and verify no results are returned:

```
#!/bin/bash

for i in $(cut -s -d: -f4 /etc/passwd | sort -u ); do
    grep -q -P "^.*?:[^:]*:$i:" /etc/group
    if [ $? -ne 0 ]; then
        echo "Group $i is referenced by /etc/passwd but does not exist in /etc/group"
    fi
done
```

Remediation:

Analyze the output of the Audit step above and perform the appropriate action to correct any discrepancies found.

References:

1. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1222, T1222.002	TA0003	M1027

6.2.4 Ensure no duplicate UIDs exist (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Although the `useradd` program will not let you create a duplicate User ID (UID), it is possible for an administrator to manually edit the `/etc/passwd` file and change the UID field.

Rationale:

Users must be assigned unique UIDs for accountability and to ensure appropriate access protections.

Audit:

Run the following script and verify no results are returned:

```
#!/bin/bash

cut -f3 -d":" /etc/passwd | sort -n | uniq -c | while read x ; do
    [ -z "$x" ] && break
    set - $x
    if [ $1 -gt 1 ]; then
        users=$(awk -F: '($3 == n) { print $1 }' n=$2 /etc/passwd | xargs)
        echo "Duplicate UID ($2): $users"
    fi
done
```

Remediation:

Based on the results of the audit script, establish unique UIDs and review all files owned by the shared UIDs to determine which UID they are supposed to belong to.

References:

1. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1078, T1078.001, T1078.003	TA0005	M1027

6.2.5 Ensure no duplicate GIDs exist (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Although the `groupadd` program will not let you create a duplicate Group ID (GID), it is possible for an administrator to manually edit the `/etc/group` file and change the GID field.

Rationale:

User groups must be assigned unique GIDs for accountability and to ensure appropriate access protections.

Audit:

Run the following script and verify no results are returned:

```
#!/bin/bash

cut -d: -f3 /etc/group | sort | uniq -d | while read x ; do
    echo "Duplicate GID ($x) in /etc/group"
done
```

Remediation:

Based on the results of the audit script, establish unique GIDs and review all files owned by the shared GID to determine which group they are supposed to belong to.

References:

1. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

Additional Information:

You can also use the `grpck` command to check for other inconsistencies in the `/etc/group` file.

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1078, T1078.001, T1078.003	TA0005	M1027

6.2.6 Ensure no duplicate user names exist (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Although the `useradd` program will not let you create a duplicate user name, it is possible for an administrator to manually edit the `/etc/passwd` file and change the username.

Rationale:

If a user is assigned a duplicate user name, it will create and have access to files with the first UID for that username in `/etc/passwd`. For example, if "test4" has a UID of 1000 and a subsequent "test4" entry has a UID of 2000, logging in as "test4" will use UID 1000. Effectively, the UID is shared, which is a security problem.

Audit:

Run the following script and verify no results are returned:

```
#!/bin/bash

cut -d: -f1 /etc/passwd | sort | uniq -d | while read -r x; do
    echo "Duplicate login name $x in /etc/passwd"
done
```

Remediation:

Based on the results of the audit script, establish unique user names for the users. File ownerships will automatically reflect the change as long as the users have unique UIDs.

References:

1. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1078, T1078.001, T1078.003	TA0004	M1027

6.2.7 Ensure no duplicate group names exist (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Although the `groupadd` program will not let you create a duplicate group name, it is possible for an administrator to manually edit the `/etc/group` file and change the group name.

Rationale:

If a group is assigned a duplicate group name, it will create and have access to files with the first GID for that group in `/etc/group`. Effectively, the GID is shared, which is a security problem.

Audit:

Run the following script and verify no results are returned:

```
#!/bin/bash

cut -d: -f1 /etc/group | sort | uniq -d | while read -r x; do
    echo "Duplicate group name $x in /etc/group"
done
```

Remediation:

Based on the results of the audit script, establish unique names for the user groups. File group ownerships will automatically reflect the change as long as the groups have unique GIDs.

References:

1. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1078, T1078.001, T1078.003	TA0004	M1027

6.2.8 Ensure root PATH Integrity (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `root` user can execute any command on the system and could be fooled into executing programs unintentionally if the `PATH` is not set correctly.

Rationale:

Including the current working directory (`.`) or other writable directory in `root`'s executable path makes it likely that an attacker can gain superuser access by forcing an administrator operating as `root` to execute a Trojan horse program.

Audit:

Run the following script and verify no results are returned:

```
#!/bin/bash

RPCV="$(sudo -Hiu root env | grep '^PATH' | cut -d= -f2)"
echo "$RPCV" | grep -q "::-" && echo "root's path contains a empty directory (::)"
echo "$RPCV" | grep -q ":$" && echo "root's path contains a trailing (:)"
for x in $(echo "$RPCV" | tr ":" " "); do
    if [ -d "$x" ]; then
        ls -ldH "$x" | awk '$9 == "." {print "PATH contains current working directory (.)."}'
        $3 != "root" {print $9, "is not owned by root"}
        substr($1,6,1) != "-" {print $9, "is group writable"}
        substr($1,9,1) != "-" {print $9, "is world writable"}'
    else
        echo "$x is not a directory"
    fi
done
```

Remediation:

Correct or justify any items discovered in the Audit step.

References:

1. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1204, T1204.002	TA0006	M1022

6.2.9 Ensure root is the only UID 0 account (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Any account with UID 0 has superuser privileges on the system.

Rationale:

This access must be limited to only the default `root` account and only from the system console. Administrative access must be through an unprivileged account using an approved mechanism as noted in Item 5.6 Ensure access to the `su` command is restricted.

Audit:

Run the following command and verify that only "root" is returned:

```
# awk -F: '($3 == 0) { print $1 }' /etc/passwd  
root
```

Remediation:

Remove any users other than `root` with UID 0 or assign them a new UID if appropriate.

References:

1. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1548, T1548.000	TA0001	M1026

6.2.10 Ensure local interactive user home directories are configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The user home directory is space defined for the particular user to set local environment variables and to store personal files. While the system administrator can establish secure permissions for users' home directories, the users can easily override these. Users can be defined in `/etc/passwd` without a home directory or with a home directory that does not actually exist.

Rationale:

Since the user is accountable for files stored in the user home directory, the user must be the owner of the directory. Group or world-writable user home directories may enable malicious users to steal or modify other users' data or to gain another user's system privileges. If the user's home directory does not exist or is unassigned, the user will be placed in `/` and will not be able to write any files or have local environment variables set.

Audit:

Run the following script to Ensure:

- local interactive user home directories exist
- Ensure local interactive users own their home directories
- Ensure local interactive user home directories are mode 750 or more restrictive

```
#!/usr/bin/env bash

{
  l_output="" l_output2="" l_heout2="" l_hoout2="" l_haout2=""
  l_valid_shells="^($( awk -F/ ' $NF != "nologin" {print}' /etc/shells | sed -rn
  '/^\/{s,/,,\|\\/,g;p}' | paste -s -d '|' - ))$"
  unset a_uarr && a_uarr=() # Clear and initialize array
  while read -r l_epu l_eph; do # Populate array with users and user home location
    a_uarr+=("$l_epu $l_eph")
  done <<< "$(awk -v pat="$l_valid_shells" -F: '$(NF) ~ pat { print $1 " " $(NF-1) }'
/etc/passwd)"
  l_asize="${#a_uarr[@]}" # Here if we want to look at number of users before
proceeding
  [ "$l_asize" -gt "10000" ] && echo -e "\n ** INFO **\n - \"$l_asize\" Local
interactive users found on the system\n - This may be a long running check\n"
  while read -r l_user l_home; do
    if [ -d "$l_home" ]; then
      l_mask='0027'
      l_max="$( printf '%o' $(( 0777 & ~$l_mask )) )"
      while read -r l_own l_mode; do
        [ "$l_user" != "$l_own" ] && l_hoout2="$l_hoout2\n - User: \"$l_user\"
Home \"$l_home\" is owned by: \"$l_own\""
        if [ $(( $l_mode & $l_mask )) -gt 0 ]; then
          l_haout2="$l_haout2\n - User: \"$l_user\" Home \"$l_home\" is mode:
\"$l_mode\" should be mode: \"$l_max\" or more restrictive"
        fi
      done <<< "$(stat -Lc '%U %a' "$l_home")"
    else
      l_heout2="$l_heout2\n - User: \"$l_user\" Home \"$l_home\" Doesn't exist"
    fi
  done <<< "$(printf '%s\n' "${a_uarr[@]}")"
  [ -z "$l_heout2" ] && l_output="$l_output\n - home directories exist" ||
l_output2="$l_output2$l_heout2"
  [ -z "$l_hoout2" ] && l_output="$l_output\n - own their home directory" ||
l_output2="$l_output2$l_hoout2"
  [ -z "$l_haout2" ] && l_output="$l_output\n - home directories are mode:
\"$l_max\" or more restrictive" || l_output2="$l_output2$l_haout2"
  [ -n "$l_output" ] && l_output="- All local interactive users:$l_output"
  if [ -z "$l_output2" ]; then # If l_output2 is empty, we pass
    echo -e "\n- Audit Result:\n ** PASS **\n - * Correctly configured *
:\n$l_output"
  else
    echo -e "\n- Audit Result:\n ** FAIL **\n - * Reasons for audit failure *
:\n$l_output2"
    [ -n "$l_output" ] && echo -e "\n- * Correctly configured * :\n$l_output"
  fi
}
```

Remediation:

If a local interactive users' home directory is undefined and/or doesn't exist, follow local site policy and perform one of the following:

- Lock the user account
- Remove the user from the system
- create a directory for the user. If undefined, edit `/etc/passwd` and add the absolute path to the directory to the last field of the user.

Run the following script to:

- Remove excessive permissions from local interactive users home directories
- Update the home directory's owner


```

#!/usr/bin/env bash

{
    l_output2=""
    l_valid_shells="^( $( awk -F\ / '$NF != "nologin" {print}' /etc/shells | sed
-rn '/^\//{s/,,\//,g;p}' | paste -s -d '|' - ) )$"
    unset a_uarr && a_uarr=() # Clear and initialize array
    while read -r l_epu l_eph; do # Populate array with users and user home
location
        a_uarr+=("$l_epu $l_eph")
        done <<< "$ (awk -v pat="$l_valid_shells" -F: '$(NF) ~ pat { print $1 " "
$(NF-1) }' /etc/passwd)"
        l_asize="${#a_uarr[@]}" # Here if we want to look at number of users
before proceeding
        [ "$l_asize" -gt "10000" ] && echo -e "\n ** INFO **\n - \"$l_asize\"
Local interactive users found on the system\n - This may be a long running
process\n"
        while read -r l_user l_home; do
            if [ -d "$l_home" ]; then
                l_mask='0027'
                l_max="$ ( printf '%o' $( ( 0777 & ~$l_mask ) ) )"
                while read -r l_own l_mode; do
                    if [ "$l_user" != "$l_own" ]; then
                        l_output2="$l_output2\n - User: \"$l_user\" Home \"$l_home\"
is owned by: \"$l_own\" \n - changing ownership to: \"$l_user\" \n"
                        chown "$l_user" "$l_home"
                    fi
                    if [ $( ( $l_mode & $l_mask ) ) -gt 0 ]; then
                        l_output2="$l_output2\n - User: \"$l_user\" Home \"$l_home\"
is mode: \"$l_mode\" should be mode: \"$l_max\" or more restrictive\n -
removing excess permissions\n"
                        chmod g-w,o-rwx "$l_home"
                    fi
                done <<< "$ (stat -Lc '%U %#a' "$l_home")"
            else
                l_output2="$l_output2\n - User: \"$l_user\" Home \"$l_home\"
Doesn't exist\n - Please create a home in accordance with local site
policy"
            fi
        done <<< "$ (printf '%s\n' "${a_uarr[@]}")"
        if [ -z "$l_output2" ]; then # If l_output2 is empty, we pass
            echo -e " - No modification needed to local interactive users home
directories"
        else
            echo -e "\n$l_output2"
        fi
    }
}

```

References:

1. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1222, T1222.002	TA0005	M1022

6.2.11 Ensure local interactive user dot files access is configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

While the system administrator can establish secure permissions for users' "dot" files, the users can easily override these.

- `.forward` file specifies an email address to forward the user's mail to.
- `.rhost` file provides the "remote authentication" database for the `rcp`, `rlogin`, and `rsh` commands and the `rcmd()` function. These files bypass the standard password-based user authentication mechanism. They specify remote hosts and users that are considered trusted (i.e. are allowed to access the local system without supplying a password)
- `.netrc` file contains data for logging into a remote host or passing authentication to an API.
- `.bash_history` file keeps track of the user's last 500 commands.

Rationale:

User configuration files with excessive or incorrect access may enable malicious users to steal or modify other users' data or to gain another user's system privileges.

Audit:

Run the following script to verify local interactive user dot files:

- Don't include `.forward`, `.rhost`, or `.netrc` files
- Are mode 0644 or more restrictive
- Are owned by the local interactive user
- Are group owned by the user's primary group
- `.bash_history` is mode 0600 or more restrictive

Note: If a `.netrc` file is required, and follows local site policy, it should have permissions of 600 or more restrictive.

```

#!/usr/bin/env bash

{
    l_output="" l_output2="" l_output3="" l_output4=""
    l_bf="" l_df="" l_nf="" l_hf=""
    l_valid_shells="^($( awk -F\ / ' $NF != "nologin" {print}' /etc/shells | sed
-rn '/^\//{s,/,\//,g;p}' | paste -s -d '|' - ))$"
    unset a_uarr && a_uarr=() # Clear and initialize array
    while read -r l_epu l_eph; do # Populate array with users and user home
location
        [[ -n "$l_epu" && -n "$l_eph" ]] && a_uarr+=("$l_epu $l_eph")
    done <<< "$ (awk -v pat="$l_valid_shells" -F: '$(NF) ~ pat { print $1 " "
$(NF-1) }' /etc/passwd)"
    l_asize="${#a_uarr[@]}" # Here if we want to look at number of users
before proceeding
    l_maxsize="1000" # Maximun number of local interactive users before
warning (Default 1,000)
    [ "$l_asize" -gt "$l_maxsize" ] && echo -e "\n  ** INFO **\n  -
\"$l_asize\" Local interactive users found on the system\n  - This may be a
long running check\n"
    file_access_chk()
    {
        l_facout2=""
        l_max="$( printf '%o' $(( 0777 & ~$l_mask )) )"
        if [ $(( $l_mode & $l_mask )) -gt 0 ]; then
            l_facout2="$l_facout2\n  - File: \"$l_hdfilename\" is mode: \"$l_mode\"
and should be mode: \"$l_max\" or more restrictive"
        fi
        if [[ ! "$l_owner" =~ ($l_user) ]]; then
            l_facout2="$l_facout2\n  - File: \"$l_hdfilename\" owned by:
\"$l_owner\" and should be owned by \"${l_user//|/ or }\""
        fi
        if [[ ! "$l_gowner" =~ ($l_group) ]]; then
            l_facout2="$l_facout2\n  - File: \"$l_hdfilename\" group owned by:
\"$l_gowner\" and should be group owned by \"${l_group//|/ or }\""
        fi
    }
    while read -r l_user l_home; do
        l_fe="" l_nout2="" l_nout3="" l_dfout2="" l_hdout2="" l_bhout2=""
        if [ -d "$l_home" ]; then
            l_group="$(id -gn "$l_user" | xargs)"
            l_group="${l_group//|/}"
            while IFS= read -r -d $'\0' l_hdfilename; do
                while read -r l_mode l_owner l_gowner; do
                    case "$(basename "$l_hdfilename")" in
                        .forward | .rhost )
                            l_fe="Y" && l_bf="Y"
                            l_dfout2="$l_dfout2\n  - File: \"$l_hdfilename\" exists" ;;
                        .netrc )
                            l_mask='0177'
                            file_access_chk
                            if [ -n "$l_facout2" ]; then
                                l_fe="Y" && l_nf="Y"
                                l_nout2="$l_facout2"
                            else
                                l_nout3="  - File: \"$l_hdfilename\" exists"
                            fi ;;
                    esac
                done
            done
        fi
    done
}

```

```

        .bash_history )
        l_mask='0177'
        file_access_chk
        if [ -n "$l_facout2" ]; then
            l_fe="Y" && l_hf="Y"
            l_bhout2="$l_facout2"
        fi ;;
    * )
        l_mask='0133'
        file_access_chk
        if [ -n "$l_facout2" ]; then
            l_fe="Y" && l_df="Y"
            l_hdout2="$l_facout2"
        fi ;;
    esac
    done <<< "$(stat -Lc '%#a %U %G' "$l_hdfilename")"
done < <(find "$l_home" -xdev -type f -name '.*' -print0)
fi
if [ "$l_fe" = "Y" ]; then
    l_output2="$l_output2\n - User: \"$l_user\" Home Directory:
\"$l_home\"
    [ -n "$l_dfout2" ] && l_output2="$l_output2$l_dfout2"
    [ -n "$l_nout2" ] && l_output2="$l_output2$l_nout2"
    [ -n "$l_bhout2" ] && l_output2="$l_output2$l_bhout2"
    [ -n "$l_hdout2" ] && l_output2="$l_output2$l_hdout2"
fi
    [ -n "$l_nout3" ] && l_output3="$l_output3\n - User: \"$l_user\" Home
Directory: \"$l_home\"$l_nout3"
    done <<< "$(printf '%s\n' "${a_uarr[@]}")"
    unset a_uarr # Remove array
    [ -n "$l_output3" ] && l_output3=" - ** Warning **\n - \".netrc\" files
should be removed unless deemed necessary\n and in accordance with local
site policy:$l_output3"
    [ -z "$l_bf" ] && l_output="$l_output\n - \".forward\" or \".rhost\"
files"
    [ -z "$l_nf" ] && l_output="$l_output\n - \".netrc\" files with
incorrect access configured"
    [ -z "$l_hf" ] && l_output="$l_output\n - \".bash_history\" files with
incorrect access configured"
    [ -z "$l_df" ] && l_output="$l_output\n - \".dot\" files with incorrect
access configured"
    [ -n "$l_output" ] && l_output=" - No local interactive users home
directories contain:$l_output"
    echo -e "$l_output4"
    if [ -z "$l_output2" ]; then # If l_output2 is empty, we pass
        echo -e "\n- Audit Result:\n ** PASS **\n - * Correctly configured *
:\n$l_output\n"
        echo -e "$l_output3\n"
    else
        echo -e "\n- Audit Result:\n ** FAIL **\n - * Reasons for audit
failure * :\n$l_output2\n"
        echo -e "$l_output3\n"
        [ -n "$l_output" ] && echo -e "- * Correctly configured *
:\n$l_output\n"
    fi
}

```

Remediation:

Making global modifications to users' files without alerting the user community can result in unexpected outages and unhappy users. Therefore, it is recommended that a monitoring policy be established to report user dot file permissions and determine the action to be taken in accordance with site policy.

The following script will:

- remove excessive permissions on `dot` files within interactive users' home directories
- change ownership of `dot` files within interactive users' home directories to the user
- change group ownership of `dot` files within interactive users' home directories to the user's primary group
- list `.forward` and `.rhost` files to be investigated and manually deleted

```

#!/usr/bin/env bash

{
    l_valid_shells="^($( awk -F\ ' $NF != "nologin" {print}' /etc/shells | sed
-rn '/^\\/{s,/,,\|/,g;p}' | paste -s -d '|' - ))$"
    unset a_uarr && a_uarr=() # Clear and initialize array
    while read -r l_epu l_eph; do # Populate array with users and user home
location
        [[ -n "$l_epu" && -n "$l_eph" ]] && a_uarr+=("$l_epu $l_eph")
        done <<< "$(awk -v pat="$l_valid_shells" -F: '$(NF) ~ pat { print $1 " "
$(NF-1) }' /etc/passwd)"
        l_asize="${#a_uarr[@]}" # Here if we want to look at number of users
before proceeding
        l_maxsize="1000" # Maximum number of local interactive users before
warning (Default 1,000)
        [ "$l_asize" -gt "$l_maxsize" ] && echo -e "\n ** INFO **\n -
\"$l_asize\" Local interactive users found on the system\n - This may be a
long running check\n"
        file_access_fix()
        {
            l_facout2=""
            l_max="$( printf '%o' $(( 0777 & ~$l_mask )) )"
            if [ $(( $l_mode & $l_mask )) -gt 0 ]; then
                echo -e " - File: \"$l_hdfilename\" is mode: \"$l_mode\" and should be
mode: \"$l_max\" or more restrictive\n - Changing to mode \"$l_max\""
                chmod "$l_chp" "$l_hdfilename"
            fi
            if [[ ! "$l_owner" =~ ($l_user) ]]; then
                echo -e " - File: \"$l_hdfilename\" owned by: \"$l_owner\" and should
be owned by \"${l_user}/// or }\" \n - Changing ownership to \"$l_user\""
                chown "$l_user" "$l_hdfilename"
            fi
            if [[ ! "$l_gowner" =~ ($l_group) ]]; then
                echo -e " - File: \"$l_hdfilename\" group owned by: \"$l_gowner\" and
should be group owned by \"${l_group}/// or }\" \n - Changing group
ownership to \"$l_group\""
                chgrp "$l_group" "$l_hdfilename"
            fi
        }
        while read -r l_user l_home; do
            if [ -d "$l_home" ]; then
                echo -e "\n - Checking user: \"$l_user\" home directory:
\"$l_home\""
                l_group="$(id -gn "$l_user" | xargs)"
                l_group="${l_group}/// or }"
                while IFS= read -r -d $'\0' l_hdfilename; do
                    while read -r l_mode l_owner l_gowner; do
                        case "$(basename "$l_hdfilename")" in
                            .forward | .rhost )
                                echo -e " - File: \"$l_hdfilename\" exists\n - Please
investigate and manually delete \"$l_hdfilename\""
                                ;;
                            .netrc )
                                l_mask='0177'
                                l_chp="u-x,go-rwx"
                                file_access_fix ;;
                            .bash_history )

```

```
        l_mask='0177'
        l_chp="u-x,go-rwx"
        file_access_fix ;;
    * )
        l_mask='0133'
        l_chp="u-x,go-wx"
        file_access_fix ;;
    esac
    done <<< "$(stat -Lc '%#a %U %G' "$l_hdfilename")"
    done < <(find "$l_home" -xdev -type f -name '.*' -print0)
fi
done <<< "$(printf '%s\n' "${a_uarr[@]}")"
unset a_uarr # Remove array
}
```

References:

- 1. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 <u>Configure Data Access Control Lists</u> Configure data access control lists based on a user's need to know. Apply data access control lists, also known as access permissions, to local and remote file systems, databases, and applications.			
v7	14.6 <u>Protect Information through Access Control Lists</u> Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.			

MITRE ATT&CK Mappings:

Techniques / Sub-techniques	Tactics	Mitigations
T1222, T1222.001, T1222.002, T1552, T1552.003, T1552.004	TA0005	M1022

Appendix: Summary Table

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1	Initial Setup		
1.1	Filesystem Configuration		
1.1.1	Configure Filesystem Kernel Modules		
1.1.1.1	Ensure mounting of squashfs filesystems is disabled (Automated)	☐	☐
1.1.1.2	Ensure mounting of udf filesystems is disabled (Automated)	☐	☐
1.1.1.3	Ensure mounting of cramfs filesystems is disabled (Automated)	☐	☐
1.1.1.4	Ensure mounting of freevxfs filesystems is disabled (Automated)	☐	☐
1.1.1.5	Ensure mounting of jffs2 filesystems is disabled (Automated)	☐	☐
1.1.1.6	Ensure mounting of hfs filesystems is disabled (Automated)	☐	☐
1.1.1.7	Ensure mounting of hfsplus filesystems is disabled (Automated)	☐	☐
1.1.2	Configure /tmp		
1.1.2.1	Ensure /tmp is a separate partition (Automated)	☐	☐
1.1.2.2	Ensure nodev option set on /tmp partition (Automated)	☐	☐
1.1.2.3	Ensure noexec option set on /tmp partition (Automated)	☐	☐
1.1.2.4	Ensure nosuid option set on /tmp partition (Automated)	☐	☐
1.1.3	Configure /var		

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.1.3.1	Ensure separate partition exists for /var (Automated)	☐	☐
1.1.3.2	Ensure nodev option set on /var partition (Automated)	☐	☐
1.1.3.3	Ensure nosuid option set on /var partition (Automated)	☐	☐
1.1.4	Configure /var/tmp		
1.1.4.1	Ensure separate partition exists for /var/tmp (Automated)	☐	☐
1.1.4.2	Ensure noexec option set on /var/tmp partition (Automated)	☐	☐
1.1.4.3	Ensure nosuid option set on /var/tmp partition (Automated)	☐	☐
1.1.4.4	Ensure nodev option set on /var/tmp partition (Automated)	☐	☐
1.1.5	Configure /var/log		
1.1.5.1	Ensure separate partition exists for /var/log (Automated)	☐	☐
1.1.5.2	Ensure nodev option set on /var/log partition (Automated)	☐	☐
1.1.5.3	Ensure noexec option set on /var/log partition (Automated)	☐	☐
1.1.5.4	Ensure nosuid option set on /var/log partition (Automated)	☐	☐
1.1.6	Configure /var/log/audit		
1.1.6.1	Ensure separate partition exists for /var/log/audit (Automated)	☐	☐
1.1.6.2	Ensure noexec option set on /var/log/audit partition (Automated)	☐	☐
1.1.6.3	Ensure nodev option set on /var/log/audit partition (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.1.6.4	Ensure nosuid option set on /var/log/audit partition (Automated)	☐	☐
1.1.7	Configure /home		
1.1.7.1	Ensure separate partition exists for /home (Automated)	☐	☐
1.1.7.2	Ensure nodev option set on /home partition (Automated)	☐	☐
1.1.7.3	Ensure nosuid option set on /home partition (Automated)	☐	☐
1.1.8	Configure /dev/shm		
1.1.8.1	Ensure /dev/shm is a separate partition (Automated)	☐	☐
1.1.8.2	Ensure nodev option set on /dev/shm partition (Automated)	☐	☐
1.1.8.3	Ensure noexec option set on /dev/shm partition (Automated)	☐	☐
1.1.8.4	Ensure nosuid option set on /dev/shm partition (Automated)	☐	☐
1.1.9	Ensure usb-storage is disabled (Automated)	☐	☐
1.2	Configure Software and Patch Management		
1.2.1	Ensure GPG keys are configured (Manual)	☐	☐
1.2.2	Ensure gpgcheck is globally activated (Automated)	☐	☐
1.2.3	Ensure package manager repositories are configured (Manual)	☐	☐
1.2.4	Ensure repo_gpgcheck is globally activated (Manual)	☐	☐
1.3	Filesystem Integrity Checking		
1.3.1	Ensure AIDE is installed (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.3.2	Ensure filesystem integrity is regularly checked (Automated)	☐	☐
1.3.3	Ensure cryptographic mechanisms are used to protect the integrity of audit tools (Automated)	☐	☐
1.4	Secure Boot Settings		
1.4.1	Ensure permissions on bootloader config are configured (Automated)	☐	☐
1.5	Additional Process Hardening		
1.5.1	Ensure address space layout randomization (ASLR) is enabled (Automated)	☐	☐
1.5.2	Ensure ptrace_scope is restricted (Automated)	☐	☐
1.5.3	Ensure core dump storage is disabled (Automated)	☐	☐
1.5.4	Ensure core dump backtraces are disabled (Automated)	☐	☐
1.6	Mandatory Access Control		
1.6.1	Configure SELinux		
1.6.1.1	Ensure SELinux is installed (Automated)	☐	☐
1.6.1.2	Ensure SELinux is not disabled in bootloader configuration (Automated)	☐	☐
1.6.1.3	Ensure SELinux policy is configured (Automated)	☐	☐
1.6.1.4	Ensure the SELinux mode is not disabled (Automated)	☐	☐
1.6.1.5	Ensure the SELinux mode is enforcing (Automated)	☐	☐
1.6.1.6	Ensure no unconfined services exist (Manual)	☐	☐
1.6.1.7	Ensure SETroubleshoot is not installed (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.6.1.8	Ensure the MCS Translation Service (mcstrans) is not installed (Automated)	☐	☐
1.7	Command Line Warning Banners		
1.7.1	Ensure message of the day is configured properly (Automated)	☐	☐
1.7.2	Ensure local login warning banner is configured properly (Automated)	☐	☐
1.7.3	Ensure remote login warning banner is configured properly (Automated)	☐	☐
1.7.4	Ensure permissions on /etc/motd are configured (Automated)	☐	☐
1.7.5	Ensure permissions on /etc/issue are configured (Automated)	☐	☐
1.7.6	Ensure permissions on /etc/issue.net are configured (Automated)	☐	☐
1.8	Ensure updates, patches, and additional security software are installed (Manual)	☐	☐
1.9	Ensure system-wide crypto policy is not legacy (Automated)	☐	☐
2	Services		
2.1	Configure Time Synchronization		
2.1.1	Ensure time synchronization is in use (Automated)	☐	☐
2.1.2	Ensure chrony is configured (Automated)	☐	☐
2.2	Configure Special Purpose Services		
2.2.1	Ensure xorg-x11-server-common is not installed (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
2.2.2	Ensure avahi is not installed (Automated)	☐	☐
2.2.3	Ensure a print server is not installed (Automated)	☐	☐
2.2.4	Ensure a dhcp server is not installed (Automated)	☐	☐
2.2.5	Ensure a dns server is not installed (Automated)	☐	☐
2.2.6	Ensure an ftp server is not installed (Automated)	☐	☐
2.2.7	Ensure a tftp server is not installed (Automated)	☐	☐
2.2.8	Ensure a web server is not installed (Automated)	☐	☐
2.2.9	Ensure IMAP and POP3 server is not installed (Automated)	☐	☐
2.2.10	Ensure Samba is not installed (Automated)	☐	☐
2.2.11	Ensure HTTP Proxy Server is not installed (Automated)	☐	☐
2.2.12	Ensure net-snmp is not installed or the snmpd service is not enabled (Automated)	☐	☐
2.2.13	Ensure telnet-server is not installed (Automated)	☐	☐
2.2.14	Ensure dnsmasq is not installed (Automated)	☐	☐
2.2.15	Ensure mail transfer agent is configured for local-only mode (Automated)	☐	☐
2.2.16	Ensure nfs-utils is not installed or the nfs-server service is masked (Automated)	☐	☐
2.2.17	Ensure rpcbind is not installed or the rpcbind services are masked (Automated)	☐	☐
2.2.18	Ensure rsync-daemon is not installed or the rsyncd service is masked (Automated)	☐	☐
2.3	Service Clients		

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
2.3.1	Ensure telnet client is not installed (Automated)	☐	☐
2.3.2	Ensure LDAP client is not installed (Automated)	☐	☐
2.3.3	Ensure FTP client is not installed (Automated)	☐	☐
2.4	Ensure nonessential services listening on the system are removed or masked (Manual)	☐	☐
3	Network Configuration		
3.1	Disable unused network protocols and devices		
3.1.1	Ensure IPv6 status is identified (Manual)	☐	☐
3.1.2	Ensure DCCP is disabled (Automated)	☐	☐
3.1.3	Ensure SCTP is disabled (Automated)	☐	☐
3.1.4	Ensure RDS is disabled (Automated)	☐	☐
3.1.5	Ensure TIPC is disabled (Automated)	☐	☐
3.2	Configure Network Parameters (Host Only)		
3.2.1	Ensure IP forwarding is disabled (Automated)	☐	☐
3.2.2	Ensure packet redirect sending is disabled (Automated)	☐	☐
3.3	Configure Network Parameters (Host and Router)		
3.3.1	Ensure source routed packets are not accepted (Automated)	☐	☐
3.3.2	Ensure ICMP redirects are not accepted (Automated)	☐	☐
3.3.3	Ensure secure ICMP redirects are not accepted (Automated)	☐	☐
3.3.4	Ensure suspicious packets are logged (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
3.3.5	Ensure broadcast ICMP requests are ignored (Automated)	☐	☐
3.3.6	Ensure bogus ICMP responses are ignored (Automated)	☐	☐
3.3.7	Ensure Reverse Path Filtering is enabled (Automated)	☐	☐
3.3.8	Ensure TCP SYN Cookies is enabled (Automated)	☐	☐
3.3.9	Ensure IPv6 router advertisements are not accepted (Automated)	☐	☐
3.4	Configure Host Based Firewall		
3.4.1	Configure a firewall utility		
3.4.1.1	Ensure nftables is installed (Automated)	☐	☐
3.4.1.2	Ensure a single firewall configuration utility is in use (Automated)	☐	☐
3.4.2	Configure firewall rules		
3.4.2.1	Ensure firewalld default zone is set (Automated)	☐	☐
3.4.2.2	Ensure at least one nftables table exists (Automated)	☐	☐
3.4.2.3	Ensure nftables base chains exist (Automated)	☐	☐
3.4.2.4	Ensure host based firewall loopback traffic is configured (Automated)	☐	☐
3.4.2.5	Ensure firewalld drops unnecessary services and ports (Manual)	☐	☐
3.4.2.6	Ensure nftables established connections are configured (Manual)	☐	☐
3.4.2.7	Ensure nftables default deny firewall policy (Automated)	☐	☐
4	Access, Authentication and Authorization		

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
4.1	Configure time-based job schedulers		
4.1.1	Ensure cron daemon is installed and enabled (Automated)	☐	☐
4.1.2	Ensure permissions on /etc/crontab are configured (Automated)	☐	☐
4.1.3	Ensure permissions on /etc/cron.hourly are configured (Automated)	☐	☐
4.1.4	Ensure permissions on /etc/cron.daily are configured (Automated)	☐	☐
4.1.5	Ensure permissions on /etc/cron.weekly are configured (Automated)	☐	☐
4.1.6	Ensure permissions on /etc/cron.monthly are configured (Automated)	☐	☐
4.1.7	Ensure permissions on /etc/cron.d are configured (Automated)	☐	☐
4.1.8	Ensure cron is restricted to authorized users (Automated)	☐	☐
4.1.9	Ensure at is restricted to authorized users (Automated)	☐	☐
4.2	Configure SSH Server		
4.2.1	Ensure permissions on /etc/ssh/sshd_config are configured (Automated)	☐	☐
4.2.2	Ensure permissions on SSH private host key files are configured (Automated)	☐	☐
4.2.3	Ensure permissions on SSH public host key files are configured (Automated)	☐	☐
4.2.4	Ensure SSH access is limited (Automated)	☐	☐
4.2.5	Ensure SSH LogLevel is appropriate (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
4.2.6	Ensure SSH PAM is enabled (Automated)	☐	☐
4.2.7	Ensure SSH root login is disabled (Automated)	☐	☐
4.2.8	Ensure SSH HostbasedAuthentication is disabled (Automated)	☐	☐
4.2.9	Ensure SSH PermitEmptyPasswords is disabled (Automated)	☐	☐
4.2.10	Ensure SSH PermitUserEnvironment is disabled (Automated)	☐	☐
4.2.11	Ensure SSH IgnoreRhosts is enabled (Automated)	☐	☐
4.2.12	Ensure SSH X11 forwarding is disabled (Automated)	☐	☐
4.2.13	Ensure SSH AllowTcpForwarding is disabled (Automated)	☐	☐
4.2.14	Ensure system-wide crypto policy is not over-ridden (Automated)	☐	☐
4.2.15	Ensure SSH warning banner is configured (Automated)	☐	☐
4.2.16	Ensure SSH MaxAuthTries is set to 4 or less (Automated)	☐	☐
4.2.17	Ensure SSH MaxStartups is configured (Automated)	☐	☐
4.2.18	Ensure SSH MaxSessions is set to 10 or less (Automated)	☐	☐
4.2.19	Ensure SSH LoginGraceTime is set to one minute or less (Automated)	☐	☐
4.2.20	Ensure SSH Idle Timeout Interval is configured (Automated)	☐	☐
4.3	Configure privilege escalation		
4.3.1	Ensure sudo is installed (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
4.3.2	Ensure sudo commands use pty (Automated)	☐	☐
4.3.3	Ensure sudo log file exists (Automated)	☐	☐
4.3.4	Ensure re-authentication for privilege escalation is not disabled globally (Automated)	☐	☐
4.3.5	Ensure sudo authentication timeout is configured correctly (Automated)	☐	☐
4.3.6	Ensure access to the su command is restricted (Automated)	☐	☐
4.4	Configure authselect		
4.4.1	Ensure custom authselect profile is used (Manual)	☐	☐
4.4.2	Ensure authselect includes with-faillock (Automated)	☐	☐
4.5	Configure PAM		
4.5.1	Ensure password creation requirements are configured (Automated)	☐	☐
4.5.2	Ensure lockout for failed password attempts is configured (Automated)	☐	☐
4.5.3	Ensure password reuse is limited (Automated)	☐	☐
4.5.4	Ensure password hashing algorithm is SHA-512 (Automated)	☐	☐
4.6	User Accounts and Environment		
4.6.1	Set Shadow Password Suite Parameters		
4.6.1.1	Ensure password expiration is 365 days or less (Automated)	☐	☐
4.6.1.2	Ensure minimum days between password changes is configured (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
4.6.1.3	Ensure password expiration warning days is 7 or more (Automated)	☐	☐
4.6.1.4	Ensure inactive password lock is 30 days or less (Automated)	☐	☐
4.6.1.5	Ensure all users last password change date is in the past (Automated)	☐	☐
4.6.2	Ensure system accounts are secured (Automated)	☐	☐
4.6.3	Ensure default user shell timeout is 900 seconds or less (Automated)	☐	☐
4.6.4	Ensure default group for the root account is GID 0 (Automated)	☐	☐
4.6.5	Ensure default user umask is 027 or more restrictive (Automated)	☐	☐
4.6.6	Ensure root password is set (Manual)	☐	☐
5	Logging and Auditing		
5.1	Configure Logging		
5.1.1	Configure rsyslog		
5.1.1.1	Ensure rsyslog is installed (Manual)	☐	☐
5.1.1.2	Ensure rsyslog service is enabled (Manual)	☐	☐
5.1.1.3	Ensure journald is configured to send logs to rsyslog (Manual)	☐	☐
5.1.1.4	Ensure rsyslog default file permissions are configured (Automated)	☐	☐
5.1.1.5	Ensure logging is configured (Manual)	☐	☐
5.1.1.6	Ensure rsyslog is configured to send logs to a remote log host (Manual)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
5.1.1.7	Ensure rsyslog is not configured to receive logs from a remote client (Automated)	☐	☐
5.1.2	Configure journald		
5.1.2.1	Ensure journald is configured to send logs to a remote log host		
5.1.2.1.1	Ensure systemd-journal-remote is installed (Manual)	☐	☐
5.1.2.1.2	Ensure systemd-journal-remote is configured (Manual)	☐	☐
5.1.2.1.3	Ensure systemd-journal-remote is enabled (Manual)	☐	☐
5.1.2.1.4	Ensure journald is not configured to receive logs from a remote client (Automated)	☐	☐
5.1.2.2	Ensure journald service is enabled (Automated)	☐	☐
5.1.2.3	Ensure journald is configured to compress large log files (Automated)	☐	☐
5.1.2.4	Ensure journald is configured to write logfiles to persistent disk (Automated)	☐	☐
5.1.2.5	Ensure journald is not configured to send logs to rsyslog (Manual)	☐	☐
5.1.2.6	Ensure journald log rotation is configured per site policy (Manual)	☐	☐
5.1.2.7	Ensure journald default file permissions configured (Manual)	☐	☐
5.1.3	Ensure all logfiles have appropriate permissions and ownership (Automated)	☐	☐
5.2	Configure System Accounting (auditd)		
5.2.1	Ensure auditing is enabled		
5.2.1.1	Ensure auditd is installed (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
5.2.1.2	Ensure auditing for processes that start prior to auditd is enabled (Automated)	☐	☐
5.2.1.3	Ensure audit_backlog_limit is sufficient (Automated)	☐	☐
5.2.1.4	Ensure auditd service is enabled (Automated)	☐	☐
5.2.2	Configure Data Retention		
5.2.2.1	Ensure audit log storage size is configured (Automated)	☐	☐
5.2.2.2	Ensure audit logs are not automatically deleted (Automated)	☐	☐
5.2.2.3	Ensure system is disabled when audit logs are full (Automated)	☐	☐
5.2.3	Configure auditd rules		
5.2.3.1	Ensure changes to system administration scope (sudoers) is collected (Automated)	☐	☐
5.2.3.2	Ensure actions as another user are always logged (Automated)	☐	☐
5.2.3.3	Ensure events that modify the sudo log file are collected (Automated)	☐	☐
5.2.3.4	Ensure events that modify date and time information are collected (Automated)	☐	☐
5.2.3.5	Ensure events that modify the system's network environment are collected (Automated)	☐	☐
5.2.3.6	Ensure use of privileged commands are collected (Automated)	☐	☐
5.2.3.7	Ensure unsuccessful file access attempts are collected (Automated)	☐	☐
5.2.3.8	Ensure events that modify user/group information are collected (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
5.2.3.9	Ensure discretionary access control permission modification events are collected (Automated)	☐	☐
5.2.3.10	Ensure successful file system mounts are collected (Automated)	☐	☐
5.2.3.11	Ensure session initiation information is collected (Automated)	☐	☐
5.2.3.12	Ensure login and logout events are collected (Automated)	☐	☐
5.2.3.13	Ensure file deletion events by users are collected (Automated)	☐	☐
5.2.3.14	Ensure events that modify the system's Mandatory Access Controls are collected (Automated)	☐	☐
5.2.3.15	Ensure successful and unsuccessful attempts to use the chcon command are recorded (Automated)	☐	☐
5.2.3.16	Ensure successful and unsuccessful attempts to use the setfacl command are recorded (Automated)	☐	☐
5.2.3.17	Ensure successful and unsuccessful attempts to use the chacl command are recorded (Automated)	☐	☐
5.2.3.18	Ensure successful and unsuccessful attempts to use the usermod command are recorded (Automated)	☐	☐
5.2.3.19	Ensure kernel module loading unloading and modification is collected (Automated)	☐	☐
5.2.3.20	Ensure the audit configuration is immutable (Automated)	☐	☐
5.2.3.21	Ensure the running and on disk configuration is the same (Manual)	☐	☐
5.2.4	Configure auditd file access		
5.2.4.1	Ensure audit log files are mode 0640 or less permissive (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
5.2.4.2	Ensure only authorized users own audit log files (Automated)	☐	☐
5.2.4.3	Ensure only authorized groups are assigned ownership of audit log files (Automated)	☐	☐
5.2.4.4	Ensure the audit log directory is 0750 or more restrictive (Automated)	☐	☐
5.2.4.5	Ensure audit configuration files are 640 or more restrictive (Automated)	☐	☐
5.2.4.6	Ensure audit configuration files are owned by root (Automated)	☐	☐
5.2.4.7	Ensure audit configuration files belong to group root (Automated)	☐	☐
5.2.4.8	Ensure audit tools are 755 or more restrictive (Automated)	☐	☐
5.2.4.9	Ensure audit tools are owned by root (Automated)	☐	☐
5.2.4.10	Ensure audit tools belong to group root (Automated)	☐	☐
5.3	Ensure logrotate is configured (Manual)	☐	☐
6	System Maintenance		
6.1	System File Permissions		
6.1.1	Ensure permissions on /etc/passwd are configured (Automated)	☐	☐
6.1.2	Ensure permissions on /etc/passwd are configured (Automated)	☐	☐
6.1.3	Ensure permissions on /etc/passwd- are configured (Automated)	☐	☐
6.1.4	Ensure permissions on /etc/group are configured (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
6.1.5	Ensure permissions on /etc/group- are configured (Automated)	☐	☐
6.1.6	Ensure permissions on /etc/shadow are configured (Automated)	☐	☐
6.1.7	Ensure permissions on /etc/shadow- are configured (Automated)	☐	☐
6.1.8	Ensure permissions on /etc/gshadow are configured (Automated)	☐	☐
6.1.9	Ensure permissions on /etc/gshadow- are configured (Automated)	☐	☐
6.1.10	Audit system file permissions (Manual)	☐	☐
6.1.11	Ensure world writable files and directories are secured (Automated)	☐	☐
6.1.12	Ensure no unowned or ungrouped files or directories exist (Automated)	☐	☐
6.1.13	Ensure SUID and SGID files are reviewed (Manual)	☐	☐
6.2	Local User and Group Settings		
6.2.1	Ensure accounts in /etc/passwd use shadowed passwords (Automated)	☐	☐
6.2.2	Ensure /etc/shadow password fields are not empty (Automated)	☐	☐
6.2.3	Ensure all groups in /etc/passwd exist in /etc/group (Automated)	☐	☐
6.2.4	Ensure no duplicate UIDs exist (Automated)	☐	☐
6.2.5	Ensure no duplicate GIDs exist (Automated)	☐	☐
6.2.6	Ensure no duplicate user names exist (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
6.2.7	Ensure no duplicate group names exist (Automated)	☐	☐
6.2.8	Ensure root PATH Integrity (Automated)	☐	☐
6.2.9	Ensure root is the only UID 0 account (Automated)	☐	☐
6.2.10	Ensure local interactive user home directories are configured (Automated)	☐	☐
6.2.11	Ensure local interactive user dot files access is configured (Automated)	☐	☐

Appendix: CIS Controls v7 IG 1 Mapped Recommendations

Recommendation		Set Correctly	
		Yes	No
1.1.2.3	Ensure noexec option set on /tmp partition	☐	☐
1.1.2.4	Ensure nosuid option set on /tmp partition	☐	☐
1.1.3.1	Ensure separate partition exists for /var	☐	☐
1.1.3.2	Ensure nodev option set on /var partition	☐	☐
1.1.3.3	Ensure nosuid option set on /var partition	☐	☐
1.1.4.1	Ensure separate partition exists for /var/tmp	☐	☐
1.1.4.2	Ensure noexec option set on /var/tmp partition	☐	☐
1.1.4.3	Ensure nosuid option set on /var/tmp partition	☐	☐
1.1.4.4	Ensure nodev option set on /var/tmp partition	☐	☐
1.1.5.2	Ensure nodev option set on /var/log partition	☐	☐
1.1.5.3	Ensure noexec option set on /var/log partition	☐	☐
1.1.5.4	Ensure nosuid option set on /var/log partition	☐	☐
1.1.6.2	Ensure noexec option set on /var/log/audit partition	☐	☐
1.1.6.3	Ensure nodev option set on /var/log/audit partition	☐	☐
1.1.6.4	Ensure nosuid option set on /var/log/audit partition	☐	☐
1.1.7.1	Ensure separate partition exists for /home	☐	☐
1.1.7.2	Ensure nodev option set on /home partition	☐	☐
1.1.7.3	Ensure nosuid option set on /home partition	☐	☐
1.1.8.2	Ensure nodev option set on /dev/shm partition	☐	☐
1.1.8.3	Ensure noexec option set on /dev/shm partition	☐	☐
1.1.8.4	Ensure nosuid option set on /dev/shm partition	☐	☐
1.2.1	Ensure GPG keys are configured	☐	☐
1.2.2	Ensure gpgcheck is globally activated	☐	☐
1.2.3	Ensure package manager repositories are configured	☐	☐
1.2.4	Ensure repo_gpgcheck is globally activated	☐	☐
1.4.1	Ensure permissions on bootloader config are configured	☐	☐
1.5.3	Ensure core dump storage is disabled	☐	☐

Recommendation		Set Correctly	
		Yes	No
1.5.4	Ensure core dump backtraces are disabled	☐	☐
1.6.1.1	Ensure SELinux is installed	☐	☐
1.6.1.2	Ensure SELinux is not disabled in bootloader configuration	☐	☐
1.6.1.3	Ensure SELinux policy is configured	☐	☐
1.6.1.4	Ensure the SELinux mode is not disabled	☐	☐
1.6.1.5	Ensure the SELinux mode is enforcing	☐	☐
1.6.1.7	Ensure SETroubleshoot is not installed	☐	☐
1.7.4	Ensure permissions on /etc/motd are configured	☐	☐
1.7.5	Ensure permissions on /etc/issue are configured	☐	☐
1.7.6	Ensure permissions on /etc/issue.net are configured	☐	☐
1.8	Ensure updates, patches, and additional security software are installed	☐	☐
2.2.13	Ensure telnet-server is not installed	☐	☐
2.3.1	Ensure telnet client is not installed	☐	☐
2.3.2	Ensure LDAP client is not installed	☐	☐
3.3.4	Ensure suspicious packets are logged	☐	☐
3.4.1.1	Ensure nftables is installed	☐	☐
3.4.1.2	Ensure a single firewall configuration utility is in use	☐	☐
3.4.2.1	Ensure firewalld default zone is set	☐	☐
3.4.2.2	Ensure at least one nftables table exists	☐	☐
3.4.2.3	Ensure nftables base chains exist	☐	☐
3.4.2.4	Ensure host based firewall loopback traffic is configured	☐	☐
3.4.2.5	Ensure firewalld drops unnecessary services and ports	☐	☐
3.4.2.6	Ensure nftables established connections are configured	☐	☐
3.4.2.7	Ensure nftables default deny firewall policy	☐	☐
4.1.2	Ensure permissions on /etc/crontab are configured	☐	☐
4.1.3	Ensure permissions on /etc/cron.hourly are configured	☐	☐
4.1.4	Ensure permissions on /etc/cron.daily are configured	☐	☐
4.1.5	Ensure permissions on /etc/cron.weekly are configured	☐	☐
4.1.6	Ensure permissions on /etc/cron.monthly are configured	☐	☐
4.1.7	Ensure permissions on /etc/cron.d are configured	☐	☐

Recommendation		Set Correctly	
		Yes	No
4.1.8	Ensure cron is restricted to authorized users	☐	☐
4.1.9	Ensure at is restricted to authorized users	☐	☐
4.2.1	Ensure permissions on /etc/ssh/sshd_config are configured	☐	☐
4.2.2	Ensure permissions on SSH private host key files are configured	☐	☐
4.2.3	Ensure permissions on SSH public host key files are configured	☐	☐
4.2.4	Ensure SSH access is limited	☐	☐
4.2.5	Ensure SSH LogLevel is appropriate	☐	☐
4.2.7	Ensure SSH root login is disabled	☐	☐
4.2.10	Ensure SSH PermitUserEnvironment is disabled	☐	☐
4.2.17	Ensure SSH MaxStartups is configured	☐	☐
4.3.1	Ensure sudo is installed	☐	☐
4.3.2	Ensure sudo commands use pty	☐	☐
4.3.4	Ensure re-authentication for privilege escalation is not disabled globally	☐	☐
4.3.5	Ensure sudo authentication timeout is configured correctly	☐	☐
4.3.6	Ensure access to the su command is restricted	☐	☐
4.6.2	Ensure system accounts are secured	☐	☐
4.6.3	Ensure default user shell timeout is 900 seconds or less	☐	☐
4.6.4	Ensure default group for the root account is GID 0	☐	☐
4.6.5	Ensure default user umask is 027 or more restrictive	☐	☐
4.6.6	Ensure root password is set	☐	☐
5.1.1.1	Ensure rsyslog is installed	☐	☐
5.1.1.2	Ensure rsyslog service is enabled	☐	☐
5.1.1.3	Ensure journald is configured to send logs to rsyslog	☐	☐
5.1.1.4	Ensure rsyslog default file permissions are configured	☐	☐
5.1.1.5	Ensure logging is configured	☐	☐
5.1.1.6	Ensure rsyslog is configured to send logs to a remote log host	☐	☐

Recommendation		Set Correctly	
		Yes	No
5.1.1.7	Ensure rsyslog is not configured to receive logs from a remote client	☐	☐
5.1.2.1.1	Ensure systemd-journal-remote is installed	☐	☐
5.1.2.1.2	Ensure systemd-journal-remote is configured	☐	☐
5.1.2.1.3	Ensure systemd-journal-remote is enabled	☐	☐
5.1.2.1.4	Ensure journald is not configured to receive logs from a remote client	☐	☐
5.1.2.2	Ensure journald service is enabled	☐	☐
5.1.2.3	Ensure journald is configured to compress large log files	☐	☐
5.1.2.4	Ensure journald is configured to write logfiles to persistent disk	☐	☐
5.1.2.5	Ensure journald is not configured to send logs to rsyslog	☐	☐
5.1.2.6	Ensure journald log rotation is configured per site policy	☐	☐
5.1.2.7	Ensure journald default file permissions configured	☐	☐
5.1.3	Ensure all logfiles have appropriate permissions and ownership	☐	☐
5.2.1.1	Ensure auditd is installed	☐	☐
5.2.1.2	Ensure auditing for processes that start prior to auditd is enabled	☐	☐
5.2.1.3	Ensure audit_backlog_limit is sufficient	☐	☐
5.2.1.4	Ensure auditd service is enabled	☐	☐
5.2.3.6	Ensure use of privileged commands are collected	☐	☐
5.2.3.12	Ensure login and logout events are collected	☐	☐
5.2.3.13	Ensure file deletion events by users are collected	☐	☐
5.2.3.15	Ensure successful and unsuccessful attempts to use the chcon command are recorded	☐	☐
5.2.3.16	Ensure successful and unsuccessful attempts to use the setfacl command are recorded	☐	☐
5.2.3.17	Ensure successful and unsuccessful attempts to use the chacl command are recorded	☐	☐
5.2.3.18	Ensure successful and unsuccessful attempts to use the usermod command are recorded	☐	☐
5.2.3.19	Ensure kernel module loading unloading and modification is collected	☐	☐

Recommendation		Set Correctly	
		Yes	No
5.2.3.20	Ensure the audit configuration is immutable	☐	☐
5.2.4.1	Ensure audit log files are mode 0640 or less permissive	☐	☐
5.2.4.2	Ensure only authorized users own audit log files	☐	☐
5.2.4.3	Ensure only authorized groups are assigned ownership of audit log files	☐	☐
5.2.4.4	Ensure the audit log directory is 0750 or more restrictive	☐	☐
5.2.4.5	Ensure audit configuration files are 640 or more restrictive	☐	☐
5.2.4.6	Ensure audit configuration files are owned by root	☐	☐
5.2.4.7	Ensure audit configuration files belong to group root	☐	☐
5.2.4.8	Ensure audit tools are 755 or more restrictive	☐	☐
5.2.4.9	Ensure audit tools are owned by root	☐	☐
5.2.4.10	Ensure audit tools belong to group root	☐	☐
6.1.10	Audit system file permissions	☐	☐
6.1.11	Ensure world writable files and directories are secured	☐	☐
6.1.13	Ensure SUID and SGID files are reviewed	☐	☐
6.2.10	Ensure local interactive user home directories are configured	☐	☐
6.2.11	Ensure local interactive user dot files access is configured	☐	☐

Appendix: CIS Controls v7 IG 2 Mapped Recommendations

Recommendation		Set Correctly	
		Yes	No
1.1.1.1	Ensure mounting of squashfs filesystems is disabled	☐	☐
1.1.1.2	Ensure mounting of udf filesystems is disabled	☐	☐
1.1.1.3	Ensure mounting of cramfs filesystems is disabled	☐	☐
1.1.1.4	Ensure mounting of freevxfs filesystems is disabled	☐	☐
1.1.1.5	Ensure mounting of jffs2 filesystems is disabled	☐	☐
1.1.1.6	Ensure mounting of hfs filesystems is disabled	☐	☐
1.1.1.7	Ensure mounting of hfsplus filesystems is disabled	☐	☐
1.1.2.1	Ensure /tmp is a separate partition	☐	☐
1.1.2.2	Ensure nodev option set on /tmp partition	☐	☐
1.1.2.3	Ensure noexec option set on /tmp partition	☐	☐
1.1.2.4	Ensure nosuid option set on /tmp partition	☐	☐
1.1.3.1	Ensure separate partition exists for /var	☐	☐
1.1.3.2	Ensure nodev option set on /var partition	☐	☐
1.1.3.3	Ensure nosuid option set on /var partition	☐	☐
1.1.4.1	Ensure separate partition exists for /var/tmp	☐	☐
1.1.4.2	Ensure noexec option set on /var/tmp partition	☐	☐
1.1.4.3	Ensure nosuid option set on /var/tmp partition	☐	☐
1.1.4.4	Ensure nodev option set on /var/tmp partition	☐	☐
1.1.5.1	Ensure separate partition exists for /var/log	☐	☐
1.1.5.2	Ensure nodev option set on /var/log partition	☐	☐
1.1.5.3	Ensure noexec option set on /var/log partition	☐	☐
1.1.5.4	Ensure nosuid option set on /var/log partition	☐	☐
1.1.6.1	Ensure separate partition exists for /var/log/audit	☐	☐
1.1.6.2	Ensure noexec option set on /var/log/audit partition	☐	☐
1.1.6.3	Ensure nodev option set on /var/log/audit partition	☐	☐
1.1.6.4	Ensure nosuid option set on /var/log/audit partition	☐	☐
1.1.7.1	Ensure separate partition exists for /home	☐	☐

Recommendation		Set Correctly	
		Yes	No
1.1.7.2	Ensure nodev option set on /home partition	☐	☐
1.1.7.3	Ensure nosuid option set on /home partition	☐	☐
1.1.8.1	Ensure /dev/shm is a separate partition	☐	☐
1.1.8.2	Ensure nodev option set on /dev/shm partition	☐	☐
1.1.8.3	Ensure noexec option set on /dev/shm partition	☐	☐
1.1.8.4	Ensure nosuid option set on /dev/shm partition	☐	☐
1.1.9	Ensure usb-storage is disabled	☐	☐
1.2.1	Ensure GPG keys are configured	☐	☐
1.2.2	Ensure gpgcheck is globally activated	☐	☐
1.2.3	Ensure package manager repositories are configured	☐	☐
1.2.4	Ensure repo_gpgcheck is globally activated	☐	☐
1.4.1	Ensure permissions on bootloader config are configured	☐	☐
1.5.1	Ensure address space layout randomization (ASLR) is enabled	☐	☐
1.5.2	Ensure ptrace_scope is restricted	☐	☐
1.5.3	Ensure core dump storage is disabled	☐	☐
1.5.4	Ensure core dump backtraces are disabled	☐	☐
1.6.1.1	Ensure SELinux is installed	☐	☐
1.6.1.2	Ensure SELinux is not disabled in bootloader configuration	☐	☐
1.6.1.3	Ensure SELinux policy is configured	☐	☐
1.6.1.4	Ensure the SELinux mode is not disabled	☐	☐
1.6.1.5	Ensure the SELinux mode is enforcing	☐	☐
1.6.1.6	Ensure no unconfined services exist	☐	☐
1.6.1.7	Ensure SETroubleshoot is not installed	☐	☐
1.6.1.8	Ensure the MCS Translation Service (mcstrans) is not installed	☐	☐
1.7.4	Ensure permissions on /etc/motd are configured	☐	☐
1.7.5	Ensure permissions on /etc/issue are configured	☐	☐
1.7.6	Ensure permissions on /etc/issue.net are configured	☐	☐
1.8	Ensure updates, patches, and additional security software are installed	☐	☐

Recommendation		Set Correctly	
		Yes	No
1.9	Ensure system-wide crypto policy is not legacy	☐	☐
2.1.1	Ensure time synchronization is in use	☐	☐
2.1.2	Ensure chrony is configured	☐	☐
2.2.1	Ensure xorg-x11-server-common is not installed	☐	☐
2.2.2	Ensure avahi is not installed	☐	☐
2.2.3	Ensure a print server is not installed	☐	☐
2.2.4	Ensure a dhcp server is not installed	☐	☐
2.2.5	Ensure a dns server is not installed	☐	☐
2.2.6	Ensure an ftp server is not installed	☐	☐
2.2.7	Ensure a tftp server is not installed	☐	☐
2.2.8	Ensure a web server is not installed	☐	☐
2.2.9	Ensure IMAP and POP3 server is not installed	☐	☐
2.2.10	Ensure Samba is not installed	☐	☐
2.2.11	Ensure HTTP Proxy Server is not installed	☐	☐
2.2.12	Ensure net-snmp is not installed or the snmpd service is not enabled	☐	☐
2.2.13	Ensure telnet-server is not installed	☐	☐
2.2.15	Ensure mail transfer agent is configured for local-only mode	☐	☐
2.2.16	Ensure nfs-utils is not installed or the nfs-server service is masked	☐	☐
2.2.17	Ensure rpcbind is not installed or the rpcbind services are masked	☐	☐
2.2.18	Ensure rsync-daemon is not installed or the rsyncd service is masked	☐	☐
2.3.1	Ensure telnet client is not installed	☐	☐
2.3.2	Ensure LDAP client is not installed	☐	☐
2.3.3	Ensure FTP client is not installed	☐	☐
2.4	Ensure nonessential services listening on the system are removed or masked	☐	☐
3.1.1	Ensure IPv6 status is identified	☐	☐
3.1.2	Ensure DCCP is disabled	☐	☐
3.1.3	Ensure SCTP is disabled	☐	☐

Recommendation		Set Correctly	
		Yes	No
3.1.4	Ensure RDS is disabled	☐	☐
3.1.5	Ensure TIPC is disabled	☐	☐
3.2.1	Ensure IP forwarding is disabled	☐	☐
3.2.2	Ensure packet redirect sending is disabled	☐	☐
3.3.1	Ensure source routed packets are not accepted	☐	☐
3.3.2	Ensure ICMP redirects are not accepted	☐	☐
3.3.3	Ensure secure ICMP redirects are not accepted	☐	☐
3.3.4	Ensure suspicious packets are logged	☐	☐
3.3.5	Ensure broadcast ICMP requests are ignored	☐	☐
3.3.6	Ensure bogus ICMP responses are ignored	☐	☐
3.3.7	Ensure Reverse Path Filtering is enabled	☐	☐
3.3.8	Ensure TCP SYN Cookies is enabled	☐	☐
3.3.9	Ensure IPv6 router advertisements are not accepted	☐	☐
3.4.1.1	Ensure nftables is installed	☐	☐
3.4.1.2	Ensure a single firewall configuration utility is in use	☐	☐
3.4.2.1	Ensure firewalld default zone is set	☐	☐
3.4.2.2	Ensure at least one nftables table exists	☐	☐
3.4.2.3	Ensure nftables base chains exist	☐	☐
3.4.2.4	Ensure host based firewall loopback traffic is configured	☐	☐
3.4.2.5	Ensure firewalld drops unnecessary services and ports	☐	☐
3.4.2.6	Ensure nftables established connections are configured	☐	☐
3.4.2.7	Ensure nftables default deny firewall policy	☐	☐
4.1.2	Ensure permissions on /etc/crontab are configured	☐	☐
4.1.3	Ensure permissions on /etc/cron.hourly are configured	☐	☐
4.1.4	Ensure permissions on /etc/cron.daily are configured	☐	☐
4.1.5	Ensure permissions on /etc/cron.weekly are configured	☐	☐
4.1.6	Ensure permissions on /etc/cron.monthly are configured	☐	☐
4.1.7	Ensure permissions on /etc/cron.d are configured	☐	☐
4.1.8	Ensure cron is restricted to authorized users	☐	☐
4.1.9	Ensure at is restricted to authorized users	☐	☐

Recommendation		Set Correctly	
		Yes	No
4.2.1	Ensure permissions on /etc/ssh/sshd_config are configured	☐	☐
4.2.2	Ensure permissions on SSH private host key files are configured	☐	☐
4.2.3	Ensure permissions on SSH public host key files are configured	☐	☐
4.2.4	Ensure SSH access is limited	☐	☐
4.2.5	Ensure SSH LogLevel is appropriate	☐	☐
4.2.7	Ensure SSH root login is disabled	☐	☐
4.2.8	Ensure SSH HostbasedAuthentication is disabled	☐	☐
4.2.9	Ensure SSH PermitEmptyPasswords is disabled	☐	☐
4.2.10	Ensure SSH PermitUserEnvironment is disabled	☐	☐
4.2.12	Ensure SSH X11 forwarding is disabled	☐	☐
4.2.13	Ensure SSH AllowTcpForwarding is disabled	☐	☐
4.2.14	Ensure system-wide crypto policy is not over-ridden	☐	☐
4.2.17	Ensure SSH MaxStartups is configured	☐	☐
4.3.1	Ensure sudo is installed	☐	☐
4.3.2	Ensure sudo commands use pty	☐	☐
4.3.3	Ensure sudo log file exists	☐	☐
4.3.4	Ensure re-authentication for privilege escalation is not disabled globally	☐	☐
4.3.5	Ensure sudo authentication timeout is configured correctly	☐	☐
4.3.6	Ensure access to the su command is restricted	☐	☐
4.4.1	Ensure custom authselect profile is used	☐	☐
4.4.2	Ensure authselect includes with-faillock	☐	☐
4.5.1	Ensure password creation requirements are configured	☐	☐
4.5.2	Ensure lockout for failed password attempts is configured	☐	☐
4.5.3	Ensure password reuse is limited	☐	☐
4.5.4	Ensure password hashing algorithm is SHA-512	☐	☐
4.6.1.1	Ensure password expiration is 365 days or less	☐	☐
4.6.1.2	Ensure minimum days between password changes is configured	☐	☐

Recommendation		Set Correctly	
		Yes	No
4.6.1.3	Ensure password expiration warning days is 7 or more	☐	☐
4.6.1.4	Ensure inactive password lock is 30 days or less	☐	☐
4.6.1.5	Ensure all users last password change date is in the past	☐	☐
4.6.2	Ensure system accounts are secured	☐	☐
4.6.3	Ensure default user shell timeout is 900 seconds or less	☐	☐
4.6.4	Ensure default group for the root account is GID 0	☐	☐
4.6.5	Ensure default user umask is 027 or more restrictive	☐	☐
4.6.6	Ensure root password is set	☐	☐
5.1.1.1	Ensure rsyslog is installed	☐	☐
5.1.1.2	Ensure rsyslog service is enabled	☐	☐
5.1.1.3	Ensure journald is configured to send logs to rsyslog	☐	☐
5.1.1.4	Ensure rsyslog default file permissions are configured	☐	☐
5.1.1.5	Ensure logging is configured	☐	☐
5.1.1.6	Ensure rsyslog is configured to send logs to a remote log host	☐	☐
5.1.1.7	Ensure rsyslog is not configured to receive logs from a remote client	☐	☐
5.1.2.1.1	Ensure systemd-journal-remote is installed	☐	☐
5.1.2.1.2	Ensure systemd-journal-remote is configured	☐	☐
5.1.2.1.3	Ensure systemd-journal-remote is enabled	☐	☐
5.1.2.1.4	Ensure journald is not configured to receive logs from a remote client	☐	☐
5.1.2.2	Ensure journald service is enabled	☐	☐
5.1.2.3	Ensure journald is configured to compress large log files	☐	☐
5.1.2.4	Ensure journald is configured to write logfiles to persistent disk	☐	☐
5.1.2.5	Ensure journald is not configured to send logs to rsyslog	☐	☐
5.1.2.6	Ensure journald log rotation is configured per site policy	☐	☐
5.1.2.7	Ensure journald default file permissions configured	☐	☐
5.1.3	Ensure all logfiles have appropriate permissions and ownership	☐	☐
5.2.1.1	Ensure auditd is installed	☐	☐

Recommendation		Set Correctly	
		Yes	No
5.2.1.2	Ensure auditing for processes that start prior to auditd is enabled	☐	☐
5.2.1.3	Ensure audit_backlog_limit is sufficient	☐	☐
5.2.1.4	Ensure auditd service is enabled	☐	☐
5.2.2.1	Ensure audit log storage size is configured	☐	☐
5.2.2.2	Ensure audit logs are not automatically deleted	☐	☐
5.2.3.1	Ensure changes to system administration scope (sudoers) is collected	☐	☐
5.2.3.2	Ensure actions as another user are always logged	☐	☐
5.2.3.3	Ensure events that modify the sudo log file are collected	☐	☐
5.2.3.4	Ensure events that modify date and time information are collected	☐	☐
5.2.3.5	Ensure events that modify the system's network environment are collected	☐	☐
5.2.3.6	Ensure use of privileged commands are collected	☐	☐
5.2.3.8	Ensure events that modify user/group information are collected	☐	☐
5.2.3.9	Ensure discretionary access control permission modification events are collected	☐	☐
5.2.3.10	Ensure successful file system mounts are collected	☐	☐
5.2.3.11	Ensure session initiation information is collected	☐	☐
5.2.3.12	Ensure login and logout events are collected	☐	☐
5.2.3.13	Ensure file deletion events by users are collected	☐	☐
5.2.3.14	Ensure events that modify the system's Mandatory Access Controls are collected	☐	☐
5.2.3.15	Ensure successful and unsuccessful attempts to use the chcon command are recorded	☐	☐
5.2.3.16	Ensure successful and unsuccessful attempts to use the setfacl command are recorded	☐	☐
5.2.3.17	Ensure successful and unsuccessful attempts to use the chacl command are recorded	☐	☐
5.2.3.18	Ensure successful and unsuccessful attempts to use the usermod command are recorded	☐	☐

Recommendation		Set Correctly	
		Yes	No
5.2.3.19	Ensure kernel module loading unloading and modification is collected	☐	☐
5.2.3.20	Ensure the audit configuration is immutable	☐	☐
5.2.3.21	Ensure the running and on disk configuration is the same	☐	☐
5.2.4.1	Ensure audit log files are mode 0640 or less permissive	☐	☐
5.2.4.2	Ensure only authorized users own audit log files	☐	☐
5.2.4.3	Ensure only authorized groups are assigned ownership of audit log files	☐	☐
5.2.4.4	Ensure the audit log directory is 0750 or more restrictive	☐	☐
5.2.4.5	Ensure audit configuration files are 640 or more restrictive	☐	☐
5.2.4.6	Ensure audit configuration files are owned by root	☐	☐
5.2.4.7	Ensure audit configuration files belong to group root	☐	☐
5.2.4.8	Ensure audit tools are 755 or more restrictive	☐	☐
5.2.4.9	Ensure audit tools are owned by root	☐	☐
5.2.4.10	Ensure audit tools belong to group root	☐	☐
5.3	Ensure logrotate is configured	☐	☐
6.1.1	Ensure permissions on /etc/passwd are configured	☐	☐
6.1.2	Ensure permissions on /etc/passwd are configured	☐	☐
6.1.3	Ensure permissions on /etc/passwd- are configured	☐	☐
6.1.4	Ensure permissions on /etc/group are configured	☐	☐
6.1.5	Ensure permissions on /etc/group- are configured	☐	☐
6.1.6	Ensure permissions on /etc/shadow are configured	☐	☐
6.1.7	Ensure permissions on /etc/shadow- are configured	☐	☐
6.1.8	Ensure permissions on /etc/gshadow are configured	☐	☐
6.1.9	Ensure permissions on /etc/gshadow- are configured	☐	☐
6.1.10	Audit system file permissions	☐	☐
6.1.11	Ensure world writable files and directories are secured	☐	☐
6.1.13	Ensure SUID and SGID files are reviewed	☐	☐
6.2.1	Ensure accounts in /etc/passwd use shadowed passwords	☐	☐
6.2.2	Ensure /etc/shadow password fields are not empty	☐	☐

Recommendation		Set Correctly	
		Yes	No
6.2.10	Ensure local interactive user home directories are configured	☐	☐
6.2.11	Ensure local interactive user dot files access is configured	☐	☐

Appendix: CIS Controls v7 IG 3 Mapped Recommendations

Recommendation		Set Correctly	
		Yes	No
1.1.1.1	Ensure mounting of squashfs filesystems is disabled	☐	☐
1.1.1.2	Ensure mounting of udf filesystems is disabled	☐	☐
1.1.1.3	Ensure mounting of cramfs filesystems is disabled	☐	☐
1.1.1.4	Ensure mounting of freevxfs filesystems is disabled	☐	☐
1.1.1.5	Ensure mounting of jffs2 filesystems is disabled	☐	☐
1.1.1.6	Ensure mounting of hfs filesystems is disabled	☐	☐
1.1.1.7	Ensure mounting of hfsplus filesystems is disabled	☐	☐
1.1.2.1	Ensure /tmp is a separate partition	☐	☐
1.1.2.2	Ensure nodev option set on /tmp partition	☐	☐
1.1.2.3	Ensure noexec option set on /tmp partition	☐	☐
1.1.2.4	Ensure nosuid option set on /tmp partition	☐	☐
1.1.3.1	Ensure separate partition exists for /var	☐	☐
1.1.3.2	Ensure nodev option set on /var partition	☐	☐
1.1.3.3	Ensure nosuid option set on /var partition	☐	☐
1.1.4.1	Ensure separate partition exists for /var/tmp	☐	☐
1.1.4.2	Ensure noexec option set on /var/tmp partition	☐	☐
1.1.4.3	Ensure nosuid option set on /var/tmp partition	☐	☐
1.1.4.4	Ensure nodev option set on /var/tmp partition	☐	☐
1.1.5.1	Ensure separate partition exists for /var/log	☐	☐
1.1.5.2	Ensure nodev option set on /var/log partition	☐	☐
1.1.5.3	Ensure noexec option set on /var/log partition	☐	☐
1.1.5.4	Ensure nosuid option set on /var/log partition	☐	☐
1.1.6.1	Ensure separate partition exists for /var/log/audit	☐	☐
1.1.6.2	Ensure noexec option set on /var/log/audit partition	☐	☐
1.1.6.3	Ensure nodev option set on /var/log/audit partition	☐	☐
1.1.6.4	Ensure nosuid option set on /var/log/audit partition	☐	☐
1.1.7.1	Ensure separate partition exists for /home	☐	☐

Recommendation		Set Correctly	
		Yes	No
1.1.7.2	Ensure nodev option set on /home partition	☐	☐
1.1.7.3	Ensure nosuid option set on /home partition	☐	☐
1.1.8.1	Ensure /dev/shm is a separate partition	☐	☐
1.1.8.2	Ensure nodev option set on /dev/shm partition	☐	☐
1.1.8.3	Ensure noexec option set on /dev/shm partition	☐	☐
1.1.8.4	Ensure nosuid option set on /dev/shm partition	☐	☐
1.1.9	Ensure usb-storage is disabled	☐	☐
1.2.1	Ensure GPG keys are configured	☐	☐
1.2.2	Ensure gpgcheck is globally activated	☐	☐
1.2.3	Ensure package manager repositories are configured	☐	☐
1.2.4	Ensure repo_gpgcheck is globally activated	☐	☐
1.3.1	Ensure AIDE is installed	☐	☐
1.3.2	Ensure filesystem integrity is regularly checked	☐	☐
1.4.1	Ensure permissions on bootloader config are configured	☐	☐
1.5.1	Ensure address space layout randomization (ASLR) is enabled	☐	☐
1.5.2	Ensure ptrace_scope is restricted	☐	☐
1.5.3	Ensure core dump storage is disabled	☐	☐
1.5.4	Ensure core dump backtraces are disabled	☐	☐
1.6.1.1	Ensure SELinux is installed	☐	☐
1.6.1.2	Ensure SELinux is not disabled in bootloader configuration	☐	☐
1.6.1.3	Ensure SELinux policy is configured	☐	☐
1.6.1.4	Ensure the SELinux mode is not disabled	☐	☐
1.6.1.5	Ensure the SELinux mode is enforcing	☐	☐
1.6.1.6	Ensure no unconfined services exist	☐	☐
1.6.1.7	Ensure SETroubleshoot is not installed	☐	☐
1.6.1.8	Ensure the MCS Translation Service (mcstrans) is not installed	☐	☐
1.7.4	Ensure permissions on /etc/motd are configured	☐	☐
1.7.5	Ensure permissions on /etc/issue are configured	☐	☐
1.7.6	Ensure permissions on /etc/issue.net are configured	☐	☐

Recommendation		Set Correctly	
		Yes	No
1.8	Ensure updates, patches, and additional security software are installed	☐	☐
1.9	Ensure system-wide crypto policy is not legacy	☐	☐
2.1.1	Ensure time synchronization is in use	☐	☐
2.1.2	Ensure chrony is configured	☐	☐
2.2.1	Ensure xorg-x11-server-common is not installed	☐	☐
2.2.2	Ensure avahi is not installed	☐	☐
2.2.3	Ensure a print server is not installed	☐	☐
2.2.4	Ensure a dhcp server is not installed	☐	☐
2.2.5	Ensure a dns server is not installed	☐	☐
2.2.6	Ensure an ftp server is not installed	☐	☐
2.2.7	Ensure a tftp server is not installed	☐	☐
2.2.8	Ensure a web server is not installed	☐	☐
2.2.9	Ensure IMAP and POP3 server is not installed	☐	☐
2.2.10	Ensure Samba is not installed	☐	☐
2.2.11	Ensure HTTP Proxy Server is not installed	☐	☐
2.2.12	Ensure net-snmp is not installed or the snmpd service is not enabled	☐	☐
2.2.13	Ensure telnet-server is not installed	☐	☐
2.2.15	Ensure mail transfer agent is configured for local-only mode	☐	☐
2.2.16	Ensure nfs-utils is not installed or the nfs-server service is masked	☐	☐
2.2.17	Ensure rpcbind is not installed or the rpcbind services are masked	☐	☐
2.2.18	Ensure rsync-daemon is not installed or the rsyncd service is masked	☐	☐
2.3.1	Ensure telnet client is not installed	☐	☐
2.3.2	Ensure LDAP client is not installed	☐	☐
2.3.3	Ensure FTP client is not installed	☐	☐
2.4	Ensure nonessential services listening on the system are removed or masked	☐	☐
3.1.1	Ensure IPv6 status is identified	☐	☐

Recommendation		Set Correctly	
		Yes	No
3.1.2	Ensure DCCP is disabled	☐	☐
3.1.3	Ensure SCTP is disabled	☐	☐
3.1.4	Ensure RDS is disabled	☐	☐
3.1.5	Ensure TIPC is disabled	☐	☐
3.2.1	Ensure IP forwarding is disabled	☐	☐
3.2.2	Ensure packet redirect sending is disabled	☐	☐
3.3.1	Ensure source routed packets are not accepted	☐	☐
3.3.2	Ensure ICMP redirects are not accepted	☐	☐
3.3.3	Ensure secure ICMP redirects are not accepted	☐	☐
3.3.4	Ensure suspicious packets are logged	☐	☐
3.3.5	Ensure broadcast ICMP requests are ignored	☐	☐
3.3.6	Ensure bogus ICMP responses are ignored	☐	☐
3.3.7	Ensure Reverse Path Filtering is enabled	☐	☐
3.3.8	Ensure TCP SYN Cookies is enabled	☐	☐
3.3.9	Ensure IPv6 router advertisements are not accepted	☐	☐
3.4.1.1	Ensure nftables is installed	☐	☐
3.4.1.2	Ensure a single firewall configuration utility is in use	☐	☐
3.4.2.1	Ensure firewalld default zone is set	☐	☐
3.4.2.2	Ensure at least one nftables table exists	☐	☐
3.4.2.3	Ensure nftables base chains exist	☐	☐
3.4.2.4	Ensure host based firewall loopback traffic is configured	☐	☐
3.4.2.5	Ensure firewalld drops unnecessary services and ports	☐	☐
3.4.2.6	Ensure nftables established connections are configured	☐	☐
3.4.2.7	Ensure nftables default deny firewall policy	☐	☐
4.1.2	Ensure permissions on /etc/crontab are configured	☐	☐
4.1.3	Ensure permissions on /etc/cron.hourly are configured	☐	☐
4.1.4	Ensure permissions on /etc/cron.daily are configured	☐	☐
4.1.5	Ensure permissions on /etc/cron.weekly are configured	☐	☐
4.1.6	Ensure permissions on /etc/cron.monthly are configured	☐	☐
4.1.7	Ensure permissions on /etc/cron.d are configured	☐	☐
4.1.8	Ensure cron is restricted to authorized users	☐	☐

Recommendation		Set Correctly	
		Yes	No
4.1.9	Ensure at is restricted to authorized users	☐	☐
4.2.1	Ensure permissions on /etc/ssh/sshd_config are configured	☐	☐
4.2.2	Ensure permissions on SSH private host key files are configured	☐	☐
4.2.3	Ensure permissions on SSH public host key files are configured	☐	☐
4.2.4	Ensure SSH access is limited	☐	☐
4.2.5	Ensure SSH LogLevel is appropriate	☐	☐
4.2.7	Ensure SSH root login is disabled	☐	☐
4.2.8	Ensure SSH HostbasedAuthentication is disabled	☐	☐
4.2.9	Ensure SSH PermitEmptyPasswords is disabled	☐	☐
4.2.10	Ensure SSH PermitUserEnvironment is disabled	☐	☐
4.2.12	Ensure SSH X11 forwarding is disabled	☐	☐
4.2.13	Ensure SSH AllowTcpForwarding is disabled	☐	☐
4.2.14	Ensure system-wide crypto policy is not over-ridden	☐	☐
4.2.16	Ensure SSH MaxAuthTries is set to 4 or less	☐	☐
4.2.17	Ensure SSH MaxStartups is configured	☐	☐
4.3.1	Ensure sudo is installed	☐	☐
4.3.2	Ensure sudo commands use pty	☐	☐
4.3.3	Ensure sudo log file exists	☐	☐
4.3.4	Ensure re-authentication for privilege escalation is not disabled globally	☐	☐
4.3.5	Ensure sudo authentication timeout is configured correctly	☐	☐
4.3.6	Ensure access to the su command is restricted	☐	☐
4.4.1	Ensure custom authselect profile is used	☐	☐
4.4.2	Ensure authselect includes with-faillock	☐	☐
4.5.1	Ensure password creation requirements are configured	☐	☐
4.5.2	Ensure lockout for failed password attempts is configured	☐	☐
4.5.3	Ensure password reuse is limited	☐	☐
4.5.4	Ensure password hashing algorithm is SHA-512	☐	☐
4.6.1.1	Ensure password expiration is 365 days or less	☐	☐

Recommendation		Set Correctly	
		Yes	No
4.6.1.2	Ensure minimum days between password changes is configured	☐	☐
4.6.1.3	Ensure password expiration warning days is 7 or more	☐	☐
4.6.1.4	Ensure inactive password lock is 30 days or less	☐	☐
4.6.1.5	Ensure all users last password change date is in the past	☐	☐
4.6.2	Ensure system accounts are secured	☐	☐
4.6.3	Ensure default user shell timeout is 900 seconds or less	☐	☐
4.6.4	Ensure default group for the root account is GID 0	☐	☐
4.6.5	Ensure default user umask is 027 or more restrictive	☐	☐
4.6.6	Ensure root password is set	☐	☐
5.1.1.1	Ensure rsyslog is installed	☐	☐
5.1.1.2	Ensure rsyslog service is enabled	☐	☐
5.1.1.3	Ensure journald is configured to send logs to rsyslog	☐	☐
5.1.1.4	Ensure rsyslog default file permissions are configured	☐	☐
5.1.1.5	Ensure logging is configured	☐	☐
5.1.1.6	Ensure rsyslog is configured to send logs to a remote log host	☐	☐
5.1.1.7	Ensure rsyslog is not configured to receive logs from a remote client	☐	☐
5.1.2.1.1	Ensure systemd-journal-remote is installed	☐	☐
5.1.2.1.2	Ensure systemd-journal-remote is configured	☐	☐
5.1.2.1.3	Ensure systemd-journal-remote is enabled	☐	☐
5.1.2.1.4	Ensure journald is not configured to receive logs from a remote client	☐	☐
5.1.2.2	Ensure journald service is enabled	☐	☐
5.1.2.3	Ensure journald is configured to compress large log files	☐	☐
5.1.2.4	Ensure journald is configured to write logfiles to persistent disk	☐	☐
5.1.2.5	Ensure journald is not configured to send logs to rsyslog	☐	☐
5.1.2.6	Ensure journald log rotation is configured per site policy	☐	☐
5.1.2.7	Ensure journald default file permissions configured	☐	☐
5.1.3	Ensure all logfiles have appropriate permissions and ownership	☐	☐

Recommendation		Set Correctly	
		Yes	No
5.2.1.1	Ensure auditd is installed	☐	☐
5.2.1.2	Ensure auditing for processes that start prior to auditd is enabled	☐	☐
5.2.1.3	Ensure audit_backlog_limit is sufficient	☐	☐
5.2.1.4	Ensure auditd service is enabled	☐	☐
5.2.2.1	Ensure audit log storage size is configured	☐	☐
5.2.2.2	Ensure audit logs are not automatically deleted	☐	☐
5.2.3.1	Ensure changes to system administration scope (sudoers) is collected	☐	☐
5.2.3.2	Ensure actions as another user are always logged	☐	☐
5.2.3.3	Ensure events that modify the sudo log file are collected	☐	☐
5.2.3.4	Ensure events that modify date and time information are collected	☐	☐
5.2.3.5	Ensure events that modify the system's network environment are collected	☐	☐
5.2.3.6	Ensure use of privileged commands are collected	☐	☐
5.2.3.7	Ensure unsuccessful file access attempts are collected	☐	☐
5.2.3.8	Ensure events that modify user/group information are collected	☐	☐
5.2.3.9	Ensure discretionary access control permission modification events are collected	☐	☐
5.2.3.10	Ensure successful file system mounts are collected	☐	☐
5.2.3.11	Ensure session initiation information is collected	☐	☐
5.2.3.12	Ensure login and logout events are collected	☐	☐
5.2.3.13	Ensure file deletion events by users are collected	☐	☐
5.2.3.14	Ensure events that modify the system's Mandatory Access Controls are collected	☐	☐
5.2.3.15	Ensure successful and unsuccessful attempts to use the chcon command are recorded	☐	☐
5.2.3.16	Ensure successful and unsuccessful attempts to use the setfacl command are recorded	☐	☐
5.2.3.17	Ensure successful and unsuccessful attempts to use the chacl command are recorded	☐	☐

Recommendation		Set Correctly	
		Yes	No
5.2.3.18	Ensure successful and unsuccessful attempts to use the usermod command are recorded	☐	☐
5.2.3.19	Ensure kernel module loading unloading and modification is collected	☐	☐
5.2.3.20	Ensure the audit configuration is immutable	☐	☐
5.2.3.21	Ensure the running and on disk configuration is the same	☐	☐
5.2.4.1	Ensure audit log files are mode 0640 or less permissive	☐	☐
5.2.4.2	Ensure only authorized users own audit log files	☐	☐
5.2.4.3	Ensure only authorized groups are assigned ownership of audit log files	☐	☐
5.2.4.4	Ensure the audit log directory is 0750 or more restrictive	☐	☐
5.2.4.5	Ensure audit configuration files are 640 or more restrictive	☐	☐
5.2.4.6	Ensure audit configuration files are owned by root	☐	☐
5.2.4.7	Ensure audit configuration files belong to group root	☐	☐
5.2.4.8	Ensure audit tools are 755 or more restrictive	☐	☐
5.2.4.9	Ensure audit tools are owned by root	☐	☐
5.2.4.10	Ensure audit tools belong to group root	☐	☐
5.3	Ensure logrotate is configured	☐	☐
6.1.1	Ensure permissions on /etc/passwd are configured	☐	☐
6.1.2	Ensure permissions on /etc/passwd are configured	☐	☐
6.1.3	Ensure permissions on /etc/passwd- are configured	☐	☐
6.1.4	Ensure permissions on /etc/group are configured	☐	☐
6.1.5	Ensure permissions on /etc/group- are configured	☐	☐
6.1.6	Ensure permissions on /etc/shadow are configured	☐	☐
6.1.7	Ensure permissions on /etc/shadow- are configured	☐	☐
6.1.8	Ensure permissions on /etc/gshadow are configured	☐	☐
6.1.9	Ensure permissions on /etc/gshadow- are configured	☐	☐
6.1.10	Audit system file permissions	☐	☐
6.1.11	Ensure world writable files and directories are secured	☐	☐
6.1.13	Ensure SUID and SGID files are reviewed	☐	☐
6.2.1	Ensure accounts in /etc/passwd use shadowed passwords	☐	☐

Recommendation		Set Correctly	
		Yes	No
6.2.2	Ensure /etc/shadow password fields are not empty	☐	☐
6.2.10	Ensure local interactive user home directories are configured	☐	☐
6.2.11	Ensure local interactive user dot files access is configured	☐	☐

Appendix: CIS Controls v7 Unmapped Recommendations

Recommendation		Set Correctly	
		Yes	No
1.3.3	Ensure cryptographic mechanisms are used to protect the integrity of audit tools	☐	☐
1.7.1	Ensure message of the day is configured properly	☐	☐
1.7.2	Ensure local login warning banner is configured properly	☐	☐
1.7.3	Ensure remote login warning banner is configured properly	☐	☐
2.2.14	Ensure dnsmasq is not installed	☐	☐
4.1.1	Ensure cron daemon is installed and enabled	☐	☐
4.2.6	Ensure SSH PAM is enabled	☐	☐
4.2.11	Ensure SSH IgnoreRhosts is enabled	☐	☐
4.2.15	Ensure SSH warning banner is configured	☐	☐
4.2.18	Ensure SSH MaxSessions is set to 10 or less	☐	☐
4.2.19	Ensure SSH LoginGraceTime is set to one minute or less	☐	☐
4.2.20	Ensure SSH Idle Timeout Interval is configured	☐	☐
5.2.2.3	Ensure system is disabled when audit logs are full	☐	☐
6.1.12	Ensure no unowned or ungrouped files or directories exist	☐	☐
6.2.3	Ensure all groups in /etc/passwd exist in /etc/group	☐	☐
6.2.4	Ensure no duplicate UIDs exist	☐	☐
6.2.5	Ensure no duplicate GIDs exist	☐	☐
6.2.6	Ensure no duplicate user names exist	☐	☐
6.2.7	Ensure no duplicate group names exist	☐	☐
6.2.8	Ensure root PATH Integrity	☐	☐
6.2.9	Ensure root is the only UID 0 account	☐	☐

Appendix: CIS Controls v8 IG 1 Mapped Recommendations

Recommendation		Set Correctly	
		Yes	No
1.1.1.4	Ensure mounting of freevxfs filesystems is disabled	☐	☐
1.1.2.3	Ensure noexec option set on /tmp partition	☐	☐
1.1.2.4	Ensure nosuid option set on /tmp partition	☐	☐
1.1.3.1	Ensure separate partition exists for /var	☐	☐
1.1.3.2	Ensure nodev option set on /var partition	☐	☐
1.1.3.3	Ensure nosuid option set on /var partition	☐	☐
1.1.4.1	Ensure separate partition exists for /var/tmp	☐	☐
1.1.4.2	Ensure noexec option set on /var/tmp partition	☐	☐
1.1.4.3	Ensure nosuid option set on /var/tmp partition	☐	☐
1.1.4.4	Ensure nodev option set on /var/tmp partition	☐	☐
1.1.5.1	Ensure separate partition exists for /var/log	☐	☐
1.1.5.2	Ensure nodev option set on /var/log partition	☐	☐
1.1.5.3	Ensure noexec option set on /var/log partition	☐	☐
1.1.5.4	Ensure nosuid option set on /var/log partition	☐	☐
1.1.6.1	Ensure separate partition exists for /var/log/audit	☐	☐
1.1.6.2	Ensure noexec option set on /var/log/audit partition	☐	☐
1.1.6.3	Ensure nodev option set on /var/log/audit partition	☐	☐
1.1.6.4	Ensure nosuid option set on /var/log/audit partition	☐	☐
1.1.7.1	Ensure separate partition exists for /home	☐	☐
1.1.7.2	Ensure nodev option set on /home partition	☐	☐
1.1.7.3	Ensure nosuid option set on /home partition	☐	☐
1.1.8.2	Ensure nodev option set on /dev/shm partition	☐	☐
1.1.8.3	Ensure noexec option set on /dev/shm partition	☐	☐
1.1.8.4	Ensure nosuid option set on /dev/shm partition	☐	☐
1.1.9	Ensure usb-storage is disabled	☐	☐
1.2.1	Ensure GPG keys are configured	☐	☐
1.2.2	Ensure gpgcheck is globally activated	☐	☐

Recommendation		Set Correctly	
		Yes	No
1.2.3	Ensure package manager repositories are configured	☐	☐
1.2.4	Ensure repo_gpgcheck is globally activated	☐	☐
1.4.1	Ensure permissions on bootloader config are configured	☐	☐
1.6.1.1	Ensure SELinux is installed	☐	☐
1.6.1.2	Ensure SELinux is not disabled in bootloader configuration	☐	☐
1.6.1.3	Ensure SELinux policy is configured	☐	☐
1.6.1.4	Ensure the SELinux mode is not disabled	☐	☐
1.6.1.5	Ensure the SELinux mode is enforcing	☐	☐
1.6.1.6	Ensure no unconfined services exist	☐	☐
1.7.4	Ensure permissions on /etc/motd are configured	☐	☐
1.7.5	Ensure permissions on /etc/issue are configured	☐	☐
1.7.6	Ensure permissions on /etc/issue.net are configured	☐	☐
1.8	Ensure updates, patches, and additional security software are installed	☐	☐
3.4.1.1	Ensure nftables is installed	☐	☐
3.4.1.2	Ensure a single firewall configuration utility is in use	☐	☐
3.4.2.1	Ensure firewalld default zone is set	☐	☐
3.4.2.2	Ensure at least one nftables table exists	☐	☐
3.4.2.3	Ensure nftables base chains exist	☐	☐
3.4.2.4	Ensure host based firewall loopback traffic is configured	☐	☐
3.4.2.5	Ensure firewalld drops unnecessary services and ports	☐	☐
3.4.2.6	Ensure nftables established connections are configured	☐	☐
3.4.2.7	Ensure nftables default deny firewall policy	☐	☐
4.1.2	Ensure permissions on /etc/crontab are configured	☐	☐
4.1.3	Ensure permissions on /etc/cron.hourly are configured	☐	☐
4.1.4	Ensure permissions on /etc/cron.daily are configured	☐	☐
4.1.5	Ensure permissions on /etc/cron.weekly are configured	☐	☐
4.1.6	Ensure permissions on /etc/cron.monthly are configured	☐	☐
4.1.7	Ensure permissions on /etc/cron.d are configured	☐	☐
4.1.8	Ensure cron is restricted to authorized users	☐	☐
4.1.9	Ensure at is restricted to authorized users	☐	☐

Recommendation		Set Correctly	
		Yes	No
4.2.1	Ensure permissions on /etc/ssh/sshd_config are configured	☐	☐
4.2.2	Ensure permissions on SSH private host key files are configured	☐	☐
4.2.3	Ensure permissions on SSH public host key files are configured	☐	☐
4.2.4	Ensure SSH access is limited	☐	☐
4.2.5	Ensure SSH LogLevel is appropriate	☐	☐
4.2.7	Ensure SSH root login is disabled	☐	☐
4.2.8	Ensure SSH HostbasedAuthentication is disabled	☐	☐
4.2.9	Ensure SSH PermitEmptyPasswords is disabled	☐	☐
4.2.10	Ensure SSH PermitUserEnvironment is disabled	☐	☐
4.2.13	Ensure SSH AllowTcpForwarding is disabled	☐	☐
4.2.17	Ensure SSH MaxStartups is configured	☐	☐
4.3.1	Ensure sudo is installed	☐	☐
4.3.2	Ensure sudo commands use pty	☐	☐
4.3.4	Ensure re-authentication for privilege escalation is not disabled globally	☐	☐
4.3.5	Ensure sudo authentication timeout is configured correctly	☐	☐
4.3.6	Ensure access to the su command is restricted	☐	☐
4.4.2	Ensure authselect includes with-faillock	☐	☐
4.5.1	Ensure password creation requirements are configured	☐	☐
4.5.2	Ensure lockout for failed password attempts is configured	☐	☐
4.5.3	Ensure password reuse is limited	☐	☐
4.6.1.1	Ensure password expiration is 365 days or less	☐	☐
4.6.1.2	Ensure minimum days between password changes is configured	☐	☐
4.6.1.3	Ensure password expiration warning days is 7 or more	☐	☐
4.6.1.4	Ensure inactive password lock is 30 days or less	☐	☐
4.6.1.5	Ensure all users last password change date is in the past	☐	☐
4.6.2	Ensure system accounts are secured	☐	☐
4.6.3	Ensure default user shell timeout is 900 seconds or less	☐	☐

Recommendation		Set Correctly	
		Yes	No
4.6.4	Ensure default group for the root account is GID 0	☐	☐
4.6.5	Ensure default user umask is 027 or more restrictive	☐	☐
4.6.6	Ensure root password is set	☐	☐
5.1.1.1	Ensure rsyslog is installed	☐	☐
5.1.1.2	Ensure rsyslog service is enabled	☐	☐
5.1.1.3	Ensure journald is configured to send logs to rsyslog	☐	☐
5.1.1.4	Ensure rsyslog default file permissions are configured	☐	☐
5.1.1.5	Ensure logging is configured	☐	☐
5.1.1.6	Ensure rsyslog is configured to send logs to a remote log host	☐	☐
5.1.1.7	Ensure rsyslog is not configured to receive logs from a remote client	☐	☐
5.1.2.1.1	Ensure systemd-journal-remote is installed	☐	☐
5.1.2.1.2	Ensure systemd-journal-remote is configured	☐	☐
5.1.2.1.3	Ensure systemd-journal-remote is enabled	☐	☐
5.1.2.1.4	Ensure journald is not configured to receive logs from a remote client	☐	☐
5.1.2.2	Ensure journald service is enabled	☐	☐
5.1.2.3	Ensure journald is configured to compress large log files	☐	☐
5.1.2.4	Ensure journald is configured to write logfiles to persistent disk	☐	☐
5.1.2.5	Ensure journald is not configured to send logs to rsyslog	☐	☐
5.1.2.6	Ensure journald log rotation is configured per site policy	☐	☐
5.1.2.7	Ensure journald default file permissions configured	☐	☐
5.1.3	Ensure all logfiles have appropriate permissions and ownership	☐	☐
5.2.1.1	Ensure auditd is installed	☐	☐
5.2.1.2	Ensure auditing for processes that start prior to auditd is enabled	☐	☐
5.2.1.3	Ensure audit_backlog_limit is sufficient	☐	☐
5.2.1.4	Ensure auditd service is enabled	☐	☐
5.2.2.1	Ensure audit log storage size is configured	☐	☐
5.2.2.2	Ensure audit logs are not automatically deleted	☐	☐

Recommendation		Set Correctly	
		Yes	No
5.2.2.3	Ensure system is disabled when audit logs are full	☐	☐
5.2.3.15	Ensure successful and unsuccessful attempts to use the chcon command are recorded	☐	☐
5.2.3.16	Ensure successful and unsuccessful attempts to use the setfacl command are recorded	☐	☐
5.2.3.17	Ensure successful and unsuccessful attempts to use the chacl command are recorded	☐	☐
5.2.3.18	Ensure successful and unsuccessful attempts to use the usermod command are recorded	☐	☐
5.2.3.20	Ensure the audit configuration is immutable	☐	☐
5.2.4.1	Ensure audit log files are mode 0640 or less permissive	☐	☐
5.2.4.2	Ensure only authorized users own audit log files	☐	☐
5.2.4.3	Ensure only authorized groups are assigned ownership of audit log files	☐	☐
5.2.4.4	Ensure the audit log directory is 0750 or more restrictive	☐	☐
5.2.4.5	Ensure audit configuration files are 640 or more restrictive	☐	☐
5.2.4.6	Ensure audit configuration files are owned by root	☐	☐
5.2.4.7	Ensure audit configuration files belong to group root	☐	☐
5.2.4.8	Ensure audit tools are 755 or more restrictive	☐	☐
5.2.4.9	Ensure audit tools are owned by root	☐	☐
5.2.4.10	Ensure audit tools belong to group root	☐	☐
5.3	Ensure logrotate is configured	☐	☐
6.1.1	Ensure permissions on /etc/passwd are configured	☐	☐
6.1.2	Ensure permissions on /etc/passwd are configured	☐	☐
6.1.3	Ensure permissions on /etc/passwd- are configured	☐	☐
6.1.4	Ensure permissions on /etc/group are configured	☐	☐
6.1.5	Ensure permissions on /etc/group- are configured	☐	☐
6.1.6	Ensure permissions on /etc/shadow are configured	☐	☐
6.1.7	Ensure permissions on /etc/shadow- are configured	☐	☐
6.1.8	Ensure permissions on /etc/gshadow are configured	☐	☐
6.1.9	Ensure permissions on /etc/gshadow- are configured	☐	☐
6.1.10	Audit system file permissions	☐	☐

Recommendation		Set Correctly	
		Yes	No
6.1.11	Ensure world writable files and directories are secured	☐	☐
6.1.12	Ensure no unowned or ungrouped files or directories exist	☐	☐
6.1.13	Ensure SUID and SGID files are reviewed	☐	☐
6.2.2	Ensure /etc/shadow password fields are not empty	☐	☐
6.2.10	Ensure local interactive user home directories are configured	☐	☐
6.2.11	Ensure local interactive user dot files access is configured	☐	☐

Appendix: CIS Controls v8 IG 2 Mapped Recommendations

Recommendation		Set Correctly	
		Yes	No
1.1.1.1	Ensure mounting of squashfs filesystems is disabled	☐	☐
1.1.1.2	Ensure mounting of udf filesystems is disabled	☐	☐
1.1.1.3	Ensure mounting of cramfs filesystems is disabled	☐	☐
1.1.1.4	Ensure mounting of freevxfs filesystems is disabled	☐	☐
1.1.1.5	Ensure mounting of jffs2 filesystems is disabled	☐	☐
1.1.1.6	Ensure mounting of hfs filesystems is disabled	☐	☐
1.1.1.7	Ensure mounting of hfsplus filesystems is disabled	☐	☐
1.1.2.1	Ensure /tmp is a separate partition	☐	☐
1.1.2.2	Ensure nodev option set on /tmp partition	☐	☐
1.1.2.3	Ensure noexec option set on /tmp partition	☐	☐
1.1.2.4	Ensure nosuid option set on /tmp partition	☐	☐
1.1.3.1	Ensure separate partition exists for /var	☐	☐
1.1.3.2	Ensure nodev option set on /var partition	☐	☐
1.1.3.3	Ensure nosuid option set on /var partition	☐	☐
1.1.4.1	Ensure separate partition exists for /var/tmp	☐	☐
1.1.4.2	Ensure noexec option set on /var/tmp partition	☐	☐
1.1.4.3	Ensure nosuid option set on /var/tmp partition	☐	☐
1.1.4.4	Ensure nodev option set on /var/tmp partition	☐	☐
1.1.5.1	Ensure separate partition exists for /var/log	☐	☐
1.1.5.2	Ensure nodev option set on /var/log partition	☐	☐
1.1.5.3	Ensure noexec option set on /var/log partition	☐	☐
1.1.5.4	Ensure nosuid option set on /var/log partition	☐	☐
1.1.6.1	Ensure separate partition exists for /var/log/audit	☐	☐
1.1.6.2	Ensure noexec option set on /var/log/audit partition	☐	☐
1.1.6.3	Ensure nodev option set on /var/log/audit partition	☐	☐
1.1.6.4	Ensure nosuid option set on /var/log/audit partition	☐	☐
1.1.7.1	Ensure separate partition exists for /home	☐	☐

Recommendation		Set Correctly	
		Yes	No
1.1.7.2	Ensure nodev option set on /home partition	☐	☐
1.1.7.3	Ensure nosuid option set on /home partition	☐	☐
1.1.8.1	Ensure /dev/shm is a separate partition	☐	☐
1.1.8.2	Ensure nodev option set on /dev/shm partition	☐	☐
1.1.8.3	Ensure noexec option set on /dev/shm partition	☐	☐
1.1.8.4	Ensure nosuid option set on /dev/shm partition	☐	☐
1.1.9	Ensure usb-storage is disabled	☐	☐
1.2.1	Ensure GPG keys are configured	☐	☐
1.2.2	Ensure gpgcheck is globally activated	☐	☐
1.2.3	Ensure package manager repositories are configured	☐	☐
1.2.4	Ensure repo_gpgcheck is globally activated	☐	☐
1.4.1	Ensure permissions on bootloader config are configured	☐	☐
1.5.1	Ensure address space layout randomization (ASLR) is enabled	☐	☐
1.5.2	Ensure ptrace_scope is restricted	☐	☐
1.6.1.1	Ensure SELinux is installed	☐	☐
1.6.1.2	Ensure SELinux is not disabled in bootloader configuration	☐	☐
1.6.1.3	Ensure SELinux policy is configured	☐	☐
1.6.1.4	Ensure the SELinux mode is not disabled	☐	☐
1.6.1.5	Ensure the SELinux mode is enforcing	☐	☐
1.6.1.6	Ensure no unconfined services exist	☐	☐
1.6.1.7	Ensure SETroubleshoot is not installed	☐	☐
1.6.1.8	Ensure the MCS Translation Service (mcstrans) is not installed	☐	☐
1.7.4	Ensure permissions on /etc/motd are configured	☐	☐
1.7.5	Ensure permissions on /etc/issue are configured	☐	☐
1.7.6	Ensure permissions on /etc/issue.net are configured	☐	☐
1.8	Ensure updates, patches, and additional security software are installed	☐	☐
1.9	Ensure system-wide crypto policy is not legacy	☐	☐
2.1.1	Ensure time synchronization is in use	☐	☐

Recommendation		Set Correctly	
		Yes	No
2.1.2	Ensure chrony is configured	☐	☐
2.2.1	Ensure xorg-x11-server-common is not installed	☐	☐
2.2.2	Ensure avahi is not installed	☐	☐
2.2.3	Ensure a print server is not installed	☐	☐
2.2.4	Ensure a dhcp server is not installed	☐	☐
2.2.5	Ensure a dns server is not installed	☐	☐
2.2.6	Ensure an ftp server is not installed	☐	☐
2.2.7	Ensure a tftp server is not installed	☐	☐
2.2.8	Ensure a web server is not installed	☐	☐
2.2.9	Ensure IMAP and POP3 server is not installed	☐	☐
2.2.10	Ensure Samba is not installed	☐	☐
2.2.11	Ensure HTTP Proxy Server is not installed	☐	☐
2.2.12	Ensure net-snmp is not installed or the snmpd service is not enabled	☐	☐
2.2.13	Ensure telnet-server is not installed	☐	☐
2.2.15	Ensure mail transfer agent is configured for local-only mode	☐	☐
2.2.16	Ensure nfs-utils is not installed or the nfs-server service is masked	☐	☐
2.2.17	Ensure rpcbind is not installed or the rpcbind services are masked	☐	☐
2.2.18	Ensure rsync-daemon is not installed or the rsyncd service is masked	☐	☐
2.3.1	Ensure telnet client is not installed	☐	☐
2.3.2	Ensure LDAP client is not installed	☐	☐
2.3.3	Ensure FTP client is not installed	☐	☐
2.4	Ensure nonessential services listening on the system are removed or masked	☐	☐
3.1.1	Ensure IPv6 status is identified	☐	☐
3.1.2	Ensure DCCP is disabled	☐	☐
3.1.3	Ensure SCTP is disabled	☐	☐
3.1.4	Ensure RDS is disabled	☐	☐
3.1.5	Ensure TIPC is disabled	☐	☐

Recommendation		Set Correctly	
		Yes	No
3.2.1	Ensure IP forwarding is disabled	☐	☐
3.2.2	Ensure packet redirect sending is disabled	☐	☐
3.3.1	Ensure source routed packets are not accepted	☐	☐
3.3.2	Ensure ICMP redirects are not accepted	☐	☐
3.3.3	Ensure secure ICMP redirects are not accepted	☐	☐
3.3.4	Ensure suspicious packets are logged	☐	☐
3.3.5	Ensure broadcast ICMP requests are ignored	☐	☐
3.3.6	Ensure bogus ICMP responses are ignored	☐	☐
3.3.7	Ensure Reverse Path Filtering is enabled	☐	☐
3.3.8	Ensure TCP SYN Cookies is enabled	☐	☐
3.3.9	Ensure IPv6 router advertisements are not accepted	☐	☐
3.4.1.1	Ensure nftables is installed	☐	☐
3.4.1.2	Ensure a single firewall configuration utility is in use	☐	☐
3.4.2.1	Ensure firewalld default zone is set	☐	☐
3.4.2.2	Ensure at least one nftables table exists	☐	☐
3.4.2.3	Ensure nftables base chains exist	☐	☐
3.4.2.4	Ensure host based firewall loopback traffic is configured	☐	☐
3.4.2.5	Ensure firewalld drops unnecessary services and ports	☐	☐
3.4.2.6	Ensure nftables established connections are configured	☐	☐
3.4.2.7	Ensure nftables default deny firewall policy	☐	☐
4.1.2	Ensure permissions on /etc/crontab are configured	☐	☐
4.1.3	Ensure permissions on /etc/cron.hourly are configured	☐	☐
4.1.4	Ensure permissions on /etc/cron.daily are configured	☐	☐
4.1.5	Ensure permissions on /etc/cron.weekly are configured	☐	☐
4.1.6	Ensure permissions on /etc/cron.monthly are configured	☐	☐
4.1.7	Ensure permissions on /etc/cron.d are configured	☐	☐
4.1.8	Ensure cron is restricted to authorized users	☐	☐
4.1.9	Ensure at is restricted to authorized users	☐	☐
4.2.1	Ensure permissions on /etc/ssh/sshd_config are configured	☐	☐
4.2.2	Ensure permissions on SSH private host key files are configured	☐	☐

Recommendation		Set Correctly	
		Yes	No
4.2.3	Ensure permissions on SSH public host key files are configured	☐	☐
4.2.4	Ensure SSH access is limited	☐	☐
4.2.5	Ensure SSH LogLevel is appropriate	☐	☐
4.2.7	Ensure SSH root login is disabled	☐	☐
4.2.8	Ensure SSH HostbasedAuthentication is disabled	☐	☐
4.2.9	Ensure SSH PermitEmptyPasswords is disabled	☐	☐
4.2.10	Ensure SSH PermitUserEnvironment is disabled	☐	☐
4.2.12	Ensure SSH X11 forwarding is disabled	☐	☐
4.2.13	Ensure SSH AllowTcpForwarding is disabled	☐	☐
4.2.14	Ensure system-wide crypto policy is not over-ridden	☐	☐
4.2.16	Ensure SSH MaxAuthTries is set to 4 or less	☐	☐
4.2.17	Ensure SSH MaxStartups is configured	☐	☐
4.3.1	Ensure sudo is installed	☐	☐
4.3.2	Ensure sudo commands use pty	☐	☐
4.3.3	Ensure sudo log file exists	☐	☐
4.3.4	Ensure re-authentication for privilege escalation is not disabled globally	☐	☐
4.3.5	Ensure sudo authentication timeout is configured correctly	☐	☐
4.3.6	Ensure access to the su command is restricted	☐	☐
4.4.1	Ensure custom authselect profile is used	☐	☐
4.4.2	Ensure authselect includes with-faillock	☐	☐
4.5.1	Ensure password creation requirements are configured	☐	☐
4.5.2	Ensure lockout for failed password attempts is configured	☐	☐
4.5.3	Ensure password reuse is limited	☐	☐
4.5.4	Ensure password hashing algorithm is SHA-512	☐	☐
4.6.1.1	Ensure password expiration is 365 days or less	☐	☐
4.6.1.2	Ensure minimum days between password changes is configured	☐	☐
4.6.1.3	Ensure password expiration warning days is 7 or more	☐	☐
4.6.1.4	Ensure inactive password lock is 30 days or less	☐	☐

Recommendation		Set Correctly	
		Yes	No
4.6.1.5	Ensure all users last password change date is in the past	☐	☐
4.6.2	Ensure system accounts are secured	☐	☐
4.6.3	Ensure default user shell timeout is 900 seconds or less	☐	☐
4.6.4	Ensure default group for the root account is GID 0	☐	☐
4.6.5	Ensure default user umask is 027 or more restrictive	☐	☐
4.6.6	Ensure root password is set	☐	☐
5.1.1.1	Ensure rsyslog is installed	☐	☐
5.1.1.2	Ensure rsyslog service is enabled	☐	☐
5.1.1.3	Ensure journald is configured to send logs to rsyslog	☐	☐
5.1.1.4	Ensure rsyslog default file permissions are configured	☐	☐
5.1.1.5	Ensure logging is configured	☐	☐
5.1.1.6	Ensure rsyslog is configured to send logs to a remote log host	☐	☐
5.1.1.7	Ensure rsyslog is not configured to receive logs from a remote client	☐	☐
5.1.2.1.1	Ensure systemd-journal-remote is installed	☐	☐
5.1.2.1.2	Ensure systemd-journal-remote is configured	☐	☐
5.1.2.1.3	Ensure systemd-journal-remote is enabled	☐	☐
5.1.2.1.4	Ensure journald is not configured to receive logs from a remote client	☐	☐
5.1.2.2	Ensure journald service is enabled	☐	☐
5.1.2.3	Ensure journald is configured to compress large log files	☐	☐
5.1.2.4	Ensure journald is configured to write logfiles to persistent disk	☐	☐
5.1.2.5	Ensure journald is not configured to send logs to rsyslog	☐	☐
5.1.2.6	Ensure journald log rotation is configured per site policy	☐	☐
5.1.2.7	Ensure journald default file permissions configured	☐	☐
5.1.3	Ensure all logfiles have appropriate permissions and ownership	☐	☐
5.2.1.1	Ensure auditd is installed	☐	☐
5.2.1.2	Ensure auditing for processes that start prior to auditd is enabled	☐	☐
5.2.1.3	Ensure audit_backlog_limit is sufficient	☐	☐

Recommendation		Set Correctly	
		Yes	No
5.2.1.4	Ensure auditd service is enabled	☐	☐
5.2.2.1	Ensure audit log storage size is configured	☐	☐
5.2.2.2	Ensure audit logs are not automatically deleted	☐	☐
5.2.2.3	Ensure system is disabled when audit logs are full	☐	☐
5.2.3.1	Ensure changes to system administration scope (sudoers) is collected	☐	☐
5.2.3.2	Ensure actions as another user are always logged	☐	☐
5.2.3.3	Ensure events that modify the sudo log file are collected	☐	☐
5.2.3.4	Ensure events that modify date and time information are collected	☐	☐
5.2.3.5	Ensure events that modify the system's network environment are collected	☐	☐
5.2.3.6	Ensure use of privileged commands are collected	☐	☐
5.2.3.7	Ensure unsuccessful file access attempts are collected	☐	☐
5.2.3.8	Ensure events that modify user/group information are collected	☐	☐
5.2.3.9	Ensure discretionary access control permission modification events are collected	☐	☐
5.2.3.10	Ensure successful file system mounts are collected	☐	☐
5.2.3.11	Ensure session initiation information is collected	☐	☐
5.2.3.12	Ensure login and logout events are collected	☐	☐
5.2.3.13	Ensure file deletion events by users are collected	☐	☐
5.2.3.14	Ensure events that modify the system's Mandatory Access Controls are collected	☐	☐
5.2.3.15	Ensure successful and unsuccessful attempts to use the chcon command are recorded	☐	☐
5.2.3.16	Ensure successful and unsuccessful attempts to use the setfacl command are recorded	☐	☐
5.2.3.17	Ensure successful and unsuccessful attempts to use the chacl command are recorded	☐	☐
5.2.3.18	Ensure successful and unsuccessful attempts to use the usermod command are recorded	☐	☐
5.2.3.19	Ensure kernel module loading unloading and modification is collected	☐	☐

Recommendation		Set Correctly	
		Yes	No
5.2.3.20	Ensure the audit configuration is immutable	☐	☐
5.2.3.21	Ensure the running and on disk configuration is the same	☐	☐
5.2.4.1	Ensure audit log files are mode 0640 or less permissive	☐	☐
5.2.4.2	Ensure only authorized users own audit log files	☐	☐
5.2.4.3	Ensure only authorized groups are assigned ownership of audit log files	☐	☐
5.2.4.4	Ensure the audit log directory is 0750 or more restrictive	☐	☐
5.2.4.5	Ensure audit configuration files are 640 or more restrictive	☐	☐
5.2.4.6	Ensure audit configuration files are owned by root	☐	☐
5.2.4.7	Ensure audit configuration files belong to group root	☐	☐
5.2.4.8	Ensure audit tools are 755 or more restrictive	☐	☐
5.2.4.9	Ensure audit tools are owned by root	☐	☐
5.2.4.10	Ensure audit tools belong to group root	☐	☐
5.3	Ensure logrotate is configured	☐	☐
6.1.1	Ensure permissions on /etc/passwd are configured	☐	☐
6.1.2	Ensure permissions on /etc/passwd are configured	☐	☐
6.1.3	Ensure permissions on /etc/passwd- are configured	☐	☐
6.1.4	Ensure permissions on /etc/group are configured	☐	☐
6.1.5	Ensure permissions on /etc/group- are configured	☐	☐
6.1.6	Ensure permissions on /etc/shadow are configured	☐	☐
6.1.7	Ensure permissions on /etc/shadow- are configured	☐	☐
6.1.8	Ensure permissions on /etc/gshadow are configured	☐	☐
6.1.9	Ensure permissions on /etc/gshadow- are configured	☐	☐
6.1.10	Audit system file permissions	☐	☐
6.1.11	Ensure world writable files and directories are secured	☐	☐
6.1.12	Ensure no unowned or ungrouped files or directories exist	☐	☐
6.1.13	Ensure SUID and SGID files are reviewed	☐	☐
6.2.1	Ensure accounts in /etc/passwd use shadowed passwords	☐	☐
6.2.2	Ensure /etc/shadow password fields are not empty	☐	☐

Recommendation		Set Correctly	
		Yes	No
6.2.10	Ensure local interactive user home directories are configured	☐	☐
6.2.11	Ensure local interactive user dot files access is configured	☐	☐

Appendix: CIS Controls v8 IG 3 Mapped Recommendations

Recommendation		Set Correctly	
		Yes	No
1.1.1.1	Ensure mounting of squashfs filesystems is disabled	☐	☐
1.1.1.2	Ensure mounting of udf filesystems is disabled	☐	☐
1.1.1.3	Ensure mounting of cramfs filesystems is disabled	☐	☐
1.1.1.4	Ensure mounting of freevxfs filesystems is disabled	☐	☐
1.1.1.5	Ensure mounting of jffs2 filesystems is disabled	☐	☐
1.1.1.6	Ensure mounting of hfs filesystems is disabled	☐	☐
1.1.1.7	Ensure mounting of hfsplus filesystems is disabled	☐	☐
1.1.2.1	Ensure /tmp is a separate partition	☐	☐
1.1.2.2	Ensure nodev option set on /tmp partition	☐	☐
1.1.2.3	Ensure noexec option set on /tmp partition	☐	☐
1.1.2.4	Ensure nosuid option set on /tmp partition	☐	☐
1.1.3.1	Ensure separate partition exists for /var	☐	☐
1.1.3.2	Ensure nodev option set on /var partition	☐	☐
1.1.3.3	Ensure nosuid option set on /var partition	☐	☐
1.1.4.1	Ensure separate partition exists for /var/tmp	☐	☐
1.1.4.2	Ensure noexec option set on /var/tmp partition	☐	☐
1.1.4.3	Ensure nosuid option set on /var/tmp partition	☐	☐
1.1.4.4	Ensure nodev option set on /var/tmp partition	☐	☐
1.1.5.1	Ensure separate partition exists for /var/log	☐	☐
1.1.5.2	Ensure nodev option set on /var/log partition	☐	☐
1.1.5.3	Ensure noexec option set on /var/log partition	☐	☐
1.1.5.4	Ensure nosuid option set on /var/log partition	☐	☐
1.1.6.1	Ensure separate partition exists for /var/log/audit	☐	☐
1.1.6.2	Ensure noexec option set on /var/log/audit partition	☐	☐
1.1.6.3	Ensure nodev option set on /var/log/audit partition	☐	☐
1.1.6.4	Ensure nosuid option set on /var/log/audit partition	☐	☐
1.1.7.1	Ensure separate partition exists for /home	☐	☐

Recommendation		Set Correctly	
		Yes	No
1.1.7.2	Ensure nodev option set on /home partition	☐	☐
1.1.7.3	Ensure nosuid option set on /home partition	☐	☐
1.1.8.1	Ensure /dev/shm is a separate partition	☐	☐
1.1.8.2	Ensure nodev option set on /dev/shm partition	☐	☐
1.1.8.3	Ensure noexec option set on /dev/shm partition	☐	☐
1.1.8.4	Ensure nosuid option set on /dev/shm partition	☐	☐
1.1.9	Ensure usb-storage is disabled	☐	☐
1.2.1	Ensure GPG keys are configured	☐	☐
1.2.2	Ensure gpgcheck is globally activated	☐	☐
1.2.3	Ensure package manager repositories are configured	☐	☐
1.2.4	Ensure repo_gpgcheck is globally activated	☐	☐
1.3.1	Ensure AIDE is installed	☐	☐
1.3.2	Ensure filesystem integrity is regularly checked	☐	☐
1.3.3	Ensure cryptographic mechanisms are used to protect the integrity of audit tools	☐	☐
1.4.1	Ensure permissions on bootloader config are configured	☐	☐
1.5.1	Ensure address space layout randomization (ASLR) is enabled	☐	☐
1.5.2	Ensure ptrace_scope is restricted	☐	☐
1.6.1.1	Ensure SELinux is installed	☐	☐
1.6.1.2	Ensure SELinux is not disabled in bootloader configuration	☐	☐
1.6.1.3	Ensure SELinux policy is configured	☐	☐
1.6.1.4	Ensure the SELinux mode is not disabled	☐	☐
1.6.1.5	Ensure the SELinux mode is enforcing	☐	☐
1.6.1.6	Ensure no unconfined services exist	☐	☐
1.6.1.7	Ensure SETroubleshoot is not installed	☐	☐
1.6.1.8	Ensure the MCS Translation Service (mcstrans) is not installed	☐	☐
1.7.4	Ensure permissions on /etc/motd are configured	☐	☐
1.7.5	Ensure permissions on /etc/issue are configured	☐	☐
1.7.6	Ensure permissions on /etc/issue.net are configured	☐	☐

Recommendation		Set Correctly	
		Yes	No
1.8	Ensure updates, patches, and additional security software are installed	☐	☐
1.9	Ensure system-wide crypto policy is not legacy	☐	☐
2.1.1	Ensure time synchronization is in use	☐	☐
2.1.2	Ensure chrony is configured	☐	☐
2.2.1	Ensure xorg-x11-server-common is not installed	☐	☐
2.2.2	Ensure avahi is not installed	☐	☐
2.2.3	Ensure a print server is not installed	☐	☐
2.2.4	Ensure a dhcp server is not installed	☐	☐
2.2.5	Ensure a dns server is not installed	☐	☐
2.2.6	Ensure an ftp server is not installed	☐	☐
2.2.7	Ensure a tftp server is not installed	☐	☐
2.2.8	Ensure a web server is not installed	☐	☐
2.2.9	Ensure IMAP and POP3 server is not installed	☐	☐
2.2.10	Ensure Samba is not installed	☐	☐
2.2.11	Ensure HTTP Proxy Server is not installed	☐	☐
2.2.12	Ensure net-snmp is not installed or the snmpd service is not enabled	☐	☐
2.2.13	Ensure telnet-server is not installed	☐	☐
2.2.15	Ensure mail transfer agent is configured for local-only mode	☐	☐
2.2.16	Ensure nfs-utils is not installed or the nfs-server service is masked	☐	☐
2.2.17	Ensure rpcbind is not installed or the rpcbind services are masked	☐	☐
2.2.18	Ensure rsync-daemon is not installed or the rsyncd service is masked	☐	☐
2.3.1	Ensure telnet client is not installed	☐	☐
2.3.2	Ensure LDAP client is not installed	☐	☐
2.3.3	Ensure FTP client is not installed	☐	☐
2.4	Ensure nonessential services listening on the system are removed or masked	☐	☐
3.1.1	Ensure IPv6 status is identified	☐	☐

Recommendation		Set Correctly	
		Yes	No
3.1.2	Ensure DCCP is disabled	☐	☐
3.1.3	Ensure SCTP is disabled	☐	☐
3.1.4	Ensure RDS is disabled	☐	☐
3.1.5	Ensure TIPC is disabled	☐	☐
3.2.1	Ensure IP forwarding is disabled	☐	☐
3.2.2	Ensure packet redirect sending is disabled	☐	☐
3.3.1	Ensure source routed packets are not accepted	☐	☐
3.3.2	Ensure ICMP redirects are not accepted	☐	☐
3.3.3	Ensure secure ICMP redirects are not accepted	☐	☐
3.3.4	Ensure suspicious packets are logged	☐	☐
3.3.5	Ensure broadcast ICMP requests are ignored	☐	☐
3.3.6	Ensure bogus ICMP responses are ignored	☐	☐
3.3.7	Ensure Reverse Path Filtering is enabled	☐	☐
3.3.8	Ensure TCP SYN Cookies is enabled	☐	☐
3.3.9	Ensure IPv6 router advertisements are not accepted	☐	☐
3.4.1.1	Ensure nftables is installed	☐	☐
3.4.1.2	Ensure a single firewall configuration utility is in use	☐	☐
3.4.2.1	Ensure firewalld default zone is set	☐	☐
3.4.2.2	Ensure at least one nftables table exists	☐	☐
3.4.2.3	Ensure nftables base chains exist	☐	☐
3.4.2.4	Ensure host based firewall loopback traffic is configured	☐	☐
3.4.2.5	Ensure firewalld drops unnecessary services and ports	☐	☐
3.4.2.6	Ensure nftables established connections are configured	☐	☐
3.4.2.7	Ensure nftables default deny firewall policy	☐	☐
4.1.2	Ensure permissions on /etc/crontab are configured	☐	☐
4.1.3	Ensure permissions on /etc/cron.hourly are configured	☐	☐
4.1.4	Ensure permissions on /etc/cron.daily are configured	☐	☐
4.1.5	Ensure permissions on /etc/cron.weekly are configured	☐	☐
4.1.6	Ensure permissions on /etc/cron.monthly are configured	☐	☐
4.1.7	Ensure permissions on /etc/cron.d are configured	☐	☐
4.1.8	Ensure cron is restricted to authorized users	☐	☐

Recommendation		Set Correctly	
		Yes	No
4.1.9	Ensure at is restricted to authorized users	☐	☐
4.2.1	Ensure permissions on /etc/ssh/sshd_config are configured	☐	☐
4.2.2	Ensure permissions on SSH private host key files are configured	☐	☐
4.2.3	Ensure permissions on SSH public host key files are configured	☐	☐
4.2.4	Ensure SSH access is limited	☐	☐
4.2.5	Ensure SSH LogLevel is appropriate	☐	☐
4.2.7	Ensure SSH root login is disabled	☐	☐
4.2.8	Ensure SSH HostbasedAuthentication is disabled	☐	☐
4.2.9	Ensure SSH PermitEmptyPasswords is disabled	☐	☐
4.2.10	Ensure SSH PermitUserEnvironment is disabled	☐	☐
4.2.12	Ensure SSH X11 forwarding is disabled	☐	☐
4.2.13	Ensure SSH AllowTcpForwarding is disabled	☐	☐
4.2.14	Ensure system-wide crypto policy is not over-ridden	☐	☐
4.2.16	Ensure SSH MaxAuthTries is set to 4 or less	☐	☐
4.2.17	Ensure SSH MaxStartups is configured	☐	☐
4.3.1	Ensure sudo is installed	☐	☐
4.3.2	Ensure sudo commands use pty	☐	☐
4.3.3	Ensure sudo log file exists	☐	☐
4.3.4	Ensure re-authentication for privilege escalation is not disabled globally	☐	☐
4.3.5	Ensure sudo authentication timeout is configured correctly	☐	☐
4.3.6	Ensure access to the su command is restricted	☐	☐
4.4.1	Ensure custom authselect profile is used	☐	☐
4.4.2	Ensure authselect includes with-faillock	☐	☐
4.5.1	Ensure password creation requirements are configured	☐	☐
4.5.2	Ensure lockout for failed password attempts is configured	☐	☐
4.5.3	Ensure password reuse is limited	☐	☐
4.5.4	Ensure password hashing algorithm is SHA-512	☐	☐
4.6.1.1	Ensure password expiration is 365 days or less	☐	☐

Recommendation		Set Correctly	
		Yes	No
4.6.1.2	Ensure minimum days between password changes is configured	☐	☐
4.6.1.3	Ensure password expiration warning days is 7 or more	☐	☐
4.6.1.4	Ensure inactive password lock is 30 days or less	☐	☐
4.6.1.5	Ensure all users last password change date is in the past	☐	☐
4.6.2	Ensure system accounts are secured	☐	☐
4.6.3	Ensure default user shell timeout is 900 seconds or less	☐	☐
4.6.4	Ensure default group for the root account is GID 0	☐	☐
4.6.5	Ensure default user umask is 027 or more restrictive	☐	☐
4.6.6	Ensure root password is set	☐	☐
5.1.1.1	Ensure rsyslog is installed	☐	☐
5.1.1.2	Ensure rsyslog service is enabled	☐	☐
5.1.1.3	Ensure journald is configured to send logs to rsyslog	☐	☐
5.1.1.4	Ensure rsyslog default file permissions are configured	☐	☐
5.1.1.5	Ensure logging is configured	☐	☐
5.1.1.6	Ensure rsyslog is configured to send logs to a remote log host	☐	☐
5.1.1.7	Ensure rsyslog is not configured to receive logs from a remote client	☐	☐
5.1.2.1.1	Ensure systemd-journal-remote is installed	☐	☐
5.1.2.1.2	Ensure systemd-journal-remote is configured	☐	☐
5.1.2.1.3	Ensure systemd-journal-remote is enabled	☐	☐
5.1.2.1.4	Ensure journald is not configured to receive logs from a remote client	☐	☐
5.1.2.2	Ensure journald service is enabled	☐	☐
5.1.2.3	Ensure journald is configured to compress large log files	☐	☐
5.1.2.4	Ensure journald is configured to write logfiles to persistent disk	☐	☐
5.1.2.5	Ensure journald is not configured to send logs to rsyslog	☐	☐
5.1.2.6	Ensure journald log rotation is configured per site policy	☐	☐
5.1.2.7	Ensure journald default file permissions configured	☐	☐
5.1.3	Ensure all logfiles have appropriate permissions and ownership	☐	☐

Recommendation		Set Correctly	
		Yes	No
5.2.1.1	Ensure auditd is installed	☐	☐
5.2.1.2	Ensure auditing for processes that start prior to auditd is enabled	☐	☐
5.2.1.3	Ensure audit_backlog_limit is sufficient	☐	☐
5.2.1.4	Ensure auditd service is enabled	☐	☐
5.2.2.1	Ensure audit log storage size is configured	☐	☐
5.2.2.2	Ensure audit logs are not automatically deleted	☐	☐
5.2.2.3	Ensure system is disabled when audit logs are full	☐	☐
5.2.3.1	Ensure changes to system administration scope (sudoers) is collected	☐	☐
5.2.3.2	Ensure actions as another user are always logged	☐	☐
5.2.3.3	Ensure events that modify the sudo log file are collected	☐	☐
5.2.3.4	Ensure events that modify date and time information are collected	☐	☐
5.2.3.5	Ensure events that modify the system's network environment are collected	☐	☐
5.2.3.6	Ensure use of privileged commands are collected	☐	☐
5.2.3.7	Ensure unsuccessful file access attempts are collected	☐	☐
5.2.3.8	Ensure events that modify user/group information are collected	☐	☐
5.2.3.9	Ensure discretionary access control permission modification events are collected	☐	☐
5.2.3.10	Ensure successful file system mounts are collected	☐	☐
5.2.3.11	Ensure session initiation information is collected	☐	☐
5.2.3.12	Ensure login and logout events are collected	☐	☐
5.2.3.13	Ensure file deletion events by users are collected	☐	☐
5.2.3.14	Ensure events that modify the system's Mandatory Access Controls are collected	☐	☐
5.2.3.15	Ensure successful and unsuccessful attempts to use the chcon command are recorded	☐	☐
5.2.3.16	Ensure successful and unsuccessful attempts to use the setfacl command are recorded	☐	☐
5.2.3.17	Ensure successful and unsuccessful attempts to use the chacl command are recorded	☐	☐

Recommendation		Set Correctly	
		Yes	No
5.2.3.18	Ensure successful and unsuccessful attempts to use the usermod command are recorded	☐	☐
5.2.3.19	Ensure kernel module loading unloading and modification is collected	☐	☐
5.2.3.20	Ensure the audit configuration is immutable	☐	☐
5.2.3.21	Ensure the running and on disk configuration is the same	☐	☐
5.2.4.1	Ensure audit log files are mode 0640 or less permissive	☐	☐
5.2.4.2	Ensure only authorized users own audit log files	☐	☐
5.2.4.3	Ensure only authorized groups are assigned ownership of audit log files	☐	☐
5.2.4.4	Ensure the audit log directory is 0750 or more restrictive	☐	☐
5.2.4.5	Ensure audit configuration files are 640 or more restrictive	☐	☐
5.2.4.6	Ensure audit configuration files are owned by root	☐	☐
5.2.4.7	Ensure audit configuration files belong to group root	☐	☐
5.2.4.8	Ensure audit tools are 755 or more restrictive	☐	☐
5.2.4.9	Ensure audit tools are owned by root	☐	☐
5.2.4.10	Ensure audit tools belong to group root	☐	☐
5.3	Ensure logrotate is configured	☐	☐
6.1.1	Ensure permissions on /etc/passwd are configured	☐	☐
6.1.2	Ensure permissions on /etc/passwd are configured	☐	☐
6.1.3	Ensure permissions on /etc/passwd- are configured	☐	☐
6.1.4	Ensure permissions on /etc/group are configured	☐	☐
6.1.5	Ensure permissions on /etc/group- are configured	☐	☐
6.1.6	Ensure permissions on /etc/shadow are configured	☐	☐
6.1.7	Ensure permissions on /etc/shadow- are configured	☐	☐
6.1.8	Ensure permissions on /etc/gshadow are configured	☐	☐
6.1.9	Ensure permissions on /etc/gshadow- are configured	☐	☐
6.1.10	Audit system file permissions	☐	☐
6.1.11	Ensure world writable files and directories are secured	☐	☐
6.1.12	Ensure no unowned or ungrouped files or directories exist	☐	☐
6.1.13	Ensure SUID and SGID files are reviewed	☐	☐

Recommendation		Set Correctly	
		Yes	No
6.2.1	Ensure accounts in <code>/etc/passwd</code> use shadowed passwords	☐	☐
6.2.2	Ensure <code>/etc/shadow</code> password fields are not empty	☐	☐
6.2.10	Ensure local interactive user home directories are configured	☐	☐
6.2.11	Ensure local interactive user dot files access is configured	☐	☐

Appendix: CIS Controls v8 Unmapped Recommendations

Recommendation		Set Correctly	
		Yes	No
1.5.3	Ensure core dump storage is disabled	☐	☐
1.5.4	Ensure core dump backtraces are disabled	☐	☐
1.7.1	Ensure message of the day is configured properly	☐	☐
1.7.2	Ensure local login warning banner is configured properly	☐	☐
1.7.3	Ensure remote login warning banner is configured properly	☐	☐
2.2.14	Ensure dnsmasq is not installed	☐	☐
4.1.1	Ensure cron daemon is installed and enabled	☐	☐
4.2.6	Ensure SSH PAM is enabled	☐	☐
4.2.11	Ensure SSH IgnoreRhosts is enabled	☐	☐
4.2.15	Ensure SSH warning banner is configured	☐	☐
4.2.18	Ensure SSH MaxSessions is set to 10 or less	☐	☐
4.2.19	Ensure SSH LoginGraceTime is set to one minute or less	☐	☐
4.2.20	Ensure SSH Idle Timeout Interval is configured	☐	☐
6.2.3	Ensure all groups in /etc/passwd exist in /etc/group	☐	☐
6.2.4	Ensure no duplicate UIDs exist	☐	☐
6.2.5	Ensure no duplicate GIDs exist	☐	☐
6.2.6	Ensure no duplicate user names exist	☐	☐
6.2.7	Ensure no duplicate group names exist	☐	☐
6.2.8	Ensure root PATH Integrity	☐	☐
6.2.9	Ensure root is the only UID 0 account	☐	☐

Appendix: Change History

Date	Version	Changes for this version
Jun 26, 2023	1.0.0	Published