

CIS RHEL8 on IBM Z Linux Benchmark

v1.0.0 - 05-27-2022

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	10
Intended Audience.....	10
Consensus Guidance	11
Typographical Conventions.....	12
Recommendation Definitions.....	13
Title.....	13
Assessment Status.....	13
Automated	13
Manual.....	13
Profile	13
Description.....	13
Rationale Statement	13
Impact Statement.....	14
Audit Procedure.....	14
Remediation Procedure.....	14
Default Value.....	14
References	14
CIS Critical Security Controls® (CIS Controls®).....	14
Additional Information.....	14
Profile Definitions	15
Acknowledgements	16
Recommendations	18
1 Initial Setup	18
1.1 Filesystem Configuration	18
1.1.1 Disable unused filesystems.....	19
1.1.1.1 Ensure mounting of cramfs filesystems is disabled (Automated)	20
1.1.1.2 Ensure mounting of vFAT filesystems is limited (Manual)	22
1.1.1.3 Ensure mounting of squashfs filesystems is disabled (Automated).....	24
1.1.1.4 Ensure mounting of udf filesystems is disabled (Automated).....	26
1.1.2 Ensure /tmp is configured (Automated)	28
1.1.3 Ensure nodev option set on /tmp partition (Automated).....	31
1.1.4 Ensure nosuid option set on /tmp partition (Automated).....	33
1.1.5 Ensure noexec option set on /tmp partition (Automated).....	35
1.1.6 Ensure separate partition exists for /var (Automated)	37
1.1.7 Ensure separate partition exists for /var/tmp (Automated).....	39

1.1.8 Ensure nodev option set on /var/tmp partition (Automated)	41
1.1.9 Ensure nosuid option set on /var/tmp partition (Automated)	42
1.1.10 Ensure noexec option set on /var/tmp partition (Automated)	43
1.1.11 Ensure separate partition exists for /var/log (Automated)	44
1.1.12 Ensure separate partition exists for /var/log/audit (Automated)	46
1.1.13 Ensure separate partition exists for /home (Automated)	48
1.1.14 Ensure nodev option set on /home partition (Automated)	50
1.1.15 Ensure nodev option set on /dev/shm partition (Automated)	52
1.1.16 Ensure nosuid option set on /dev/shm partition (Automated)	54
1.1.17 Ensure noexec option set on /dev/shm partition (Automated)	56
1.1.18 Ensure nodev option set on removable media partitions (Manual)	58
1.1.19 Ensure nosuid option set on removable media partitions (Manual)	59
1.1.20 Ensure noexec option set on removable media partitions (Manual)	60
1.1.21 Ensure sticky bit is set on all world-writable directories (Automated)	61
1.1.22 Disable Automounting (Automated)	63
1.1.23 Disable USB Storage (Automated)	65
1.1.24 Mount devices read-only that do not require write access (Optional) (Manual)	67
1.2 Configure Software Updates	69
1.2.1 Ensure GPG keys are configured (Manual)	70
1.2.2 Ensure gpgcheck is globally activated (Automated)	71
1.2.3 Ensure package manager repositories are configured (Manual)	73
1.3 Configure sudo	74
1.3.1 Ensure sudo is installed (Automated)	75
1.3.2 Ensure sudo commands use pty (Automated)	77
1.3.3 Ensure sudo log file exists (Automated)	79
1.3.4 Ensure sudo-specific activity logging (Manual)	81
1.4 Filesystem Integrity Checking	83
1.4.1 Ensure AIDE is installed (Automated)	84
1.4.2 Ensure filesystem integrity is regularly checked (Automated)	86
1.5 Secure Boot Settings	88
1.5.1 Ensure a secure IPL configuration (Manual)	89
1.5.2 Ensure permissions on the boot loader configuration are configured (Automated)	90
1.5.3 Ensure boot loader password is set (Automated)	92
1.5.4 Ensure authentication required for single user mode (Automated)	94
1.5.5 Enable secure boot as IPL parameter (Manual)	95
1.6 Additional Process Hardening	96
1.6.1 Ensure core dumps are restricted (Automated)	97
1.6.2 Ensure address space layout randomization (ASLR) is enabled (Automated)	99
1.7 Mandatory Access Control	101
1.7.1 Configure SELinux	102
1.7.1.1 Ensure SELinux is installed (Automated)	104
1.7.1.2 Ensure SELinux is not disabled in boot loader configuration (Automated)	105
1.7.1.3 Ensure SELinux policy is configured (Automated)	107
1.7.1.4 Ensure the SELinux state is set to 'enforcing' (Automated)	109
1.7.1.5 Ensure no unconfined services exist (Automated)	111
1.7.1.6 Ensure SETroubleshoot is not installed (Automated)	113
1.7.1.7 Ensure the MCS Translation Service (mcstrans) is not installed (Automated)	114
1.7.1.8 Ensure SELinux is not disabled in zipl boot loader (Manual)	115
1.8 Warning Banners	117
1.8.1 Command Line Warning Banners	118
1.8.1.1 Ensure message of the day is configured properly (Automated)	119
1.8.1.2 Ensure local login warning banner is configured properly (Automated)	120

1.8.1.3 Ensure remote login warning banner is configured properly (Automated)	121
1.8.1.4 Ensure permissions on /etc/motd are configured (Automated)	122
1.8.1.5 Ensure permissions on /etc/issue are configured (Automated)	123
1.8.1.6 Ensure permissions on /etc/issue.net are configured (Automated)	124
1.8.2 Ensure GDM login banner is configured (Automated)	125
1.9 Ensure updates, patches, and additional security software are installed (Manual)	126
1.10 Ensure system-wide crypto policy is not legacy (Automated)	128
1.11 Ensure system-wide crypto policy is FUTURE or FIPS (Automated)	130
1.12 Ensure that the kernel has loaded all s390-specific crypto modules (Manual)	132
2 Services.....	134
2.1 inetd Services	135
2.1.1 Ensure xinetd is not installed (Automated)	136
2.2 Special Purpose Services	137
2.2.1 Time Synchronization	138
2.2.1.1 Ensure time synchronization is in use (Manual)	139
2.2.1.2 Ensure chrony is configured (Automated).....	141
2.2.2 Ensure X Window System is not installed (Automated).....	143
2.2.3 Ensure rsync service is not enabled (Automated)	145
2.2.4 Ensure Avahi Server is not enabled (Automated)	147
2.2.5 Ensure SNMP Server is not enabled (Automated).....	149
2.2.6 Disable the pcnfsd service (Automated)	151
2.2.7 Disable the tcpspray service (Automated)	153
2.2.8 Ensure HTTP Proxy Server is not enabled (Automated)	155
2.2.9 Ensure Samba is not enabled (Automated)	157
2.2.10 Ensure IMAP and POP3 server is not enabled (Automated)	159
2.2.11 Ensure HTTP server is not enabled (Automated).....	161
2.2.12 Ensure FTP Server is not enabled (Automated)	163
2.2.13 Ensure DNS Server is not enabled (Automated).....	165
2.2.14 Ensure NFS is not enabled (Automated)	167
2.2.15 Ensure RPC is not enabled (Automated).....	169
2.2.16 Ensure LDAP server is not enabled (Automated).....	171
2.2.17 Ensure DHCP Server is not enabled (Automated)	173
2.2.18 Ensure CUPS is not enabled (Automated)	175
2.2.19 Ensure NIS Server is not enabled (Automated)	177
2.2.20 Ensure mail transfer agent is configured for local-only mode (Automated).....	179
2.3 Service Clients.....	181
2.3.1 Ensure NIS Client is not installed (Automated).....	182
2.3.2 Ensure telnet client is not installed (Automated)	184
2.3.3 Ensure LDAP client is not installed (Automated).....	186
3 Network Configuration	188
3.1 Network Parameters (Host Only).....	189
3.1.1 Ensure IP forwarding is disabled (Automated).....	190
3.1.2 Ensure packet redirect sending is disabled (Automated).....	192
3.1.3 Enable SMC-D pre-load libraries (Manual).....	193
3.1.4 Enable VEPA mode traffic for OSA network devices (Manual)	194
3.1.5 Enable traffic isolation between OSA cards (Manual)	195
3.2 Network Parameters (Host and Router).....	196
3.2.1 Ensure source routed packets are not accepted (Automated)	197
3.2.2 Ensure ICMP redirects are not accepted (Automated)	199
3.2.3 Ensure secure ICMP redirects are not accepted (Automated)	201
3.2.4 Ensure suspicious packets are logged (Automated)	202

3.2.5 Ensure broadcast ICMP requests are ignored (Automated)	204
3.2.6 Ensure bogus ICMP responses are ignored (Automated)	205
3.2.7 Ensure reverse path filtering is enabled (Automated).....	206
3.2.8 Ensure TCP SYN Cookies is enabled (Automated)	208
3.2.9 Ensure IPv6 router advertisements are not accepted (Automated).....	209
3.3 Uncommon Network Protocols	210
3.3.1 Ensure DCCP is disabled (Automated)	211
3.3.2 Disable forwarding of packets (Automated)	213
3.3.3 Ensure SCTP is disabled (Automated).....	215
3.3.4 Ensure RDS is disabled (Automated).....	217
3.3.5 Ensure TIPC is disabled (Automated)	219
3.4 Firewall Configuration	220
3.4.1 Ensure Firewall software is installed.....	221
3.4.1.1 Ensure a Firewall package is installed (Automated).....	222
3.4.2 Configure firewalld	224
3.4.2.1 Ensure firewalld service is enabled and running (Automated)	225
3.4.2.2 Ensure iptables is not enabled (Automated)	227
3.4.2.3 Ensure nftables is not enabled (Automated).....	229
3.4.2.4 Ensure default zone is set (Automated).....	231
3.4.2.5 Ensure network interfaces are assigned to appropriate zone (Manual)	233
3.4.2.6 Ensure unnecessary services and ports are not accepted (Manual)	235
3.4.3 Configure nftables.....	237
3.4.3.1 Ensure iptables are flushed (Manual).....	240
3.4.3.2 Ensure a table exists (Automated)	242
3.4.3.3 Ensure base chains exist (Automated).....	244
3.4.3.4 Ensure loopback traffic is configured (Automated)	246
3.4.3.5 Ensure outbound and established connections are configured (Manual)	248
3.4.3.6 Ensure default deny firewall policy (Automated)	250
3.4.3.7 Ensure nftables service is enabled (Automated)	252
3.4.3.8 Ensure nftables rules are permanent (Automated).....	253
3.4.4 Configure iptables.....	256
3.4.4.1 Configure software	257
3.4.4.1.1 Ensure iptables packages are installed (Automated).....	258
3.4.4.1.2 Ensure nftables is not installed (Automated).....	259
3.4.4.1.3 Ensure firewalld is not installed or stopped and masked (Automated).....	260
3.4.4.2 Configure IPv4 iptables	262
3.4.4.2.1 Ensure default deny firewall policy (Automated)	263
3.4.4.2.2 Ensure loopback traffic is configured (Automated)	265
3.4.4.2.3 Ensure outbound and established connections are configured (Manual).....	267
3.4.4.2.4 Ensure firewall rules exist for all open ports (Automated)	269
3.4.4.3 Configure IPv6 ip6tables	271
3.4.4.3.1 Ensure IPv6 default deny firewall policy (Automated).....	272
3.4.4.3.2 Ensure IPv6 loopback traffic is configured (Automated)	274
3.4.4.3.3 Ensure IPv6 outbound and established connections are configured (Manual)	276
3.4.4.3.4 Ensure IPv6 firewall rules exist for all open ports (Manual).....	278
3.5 Ensure wireless interfaces are disabled (Automated)	280
3.6 Disable IPv6 (Manual).....	282
3.7 Choose hardware-accelerated cipher suites for applications using SSL/TLS (Manual).....	284
4 Logging and Auditing.....	285
4.1 Configure System Accounting (auditd)	286
4.1.1 Ensure auditing is enabled	287

4.1.1.1 Ensure auditd is installed (Automated).....	288
4.1.1.2 Ensure auditd service is enabled (Automated).....	289
4.1.1.3 Enable auditing for processes that start prior to auditd (Automated).....	291
4.1.1.4 Extend the audit backlog limit for the audit daemon (Automated)	293
4.1.2 Configure Data Retention.....	295
4.1.2.1 Ensure audit log storage size is configured (Automated).....	296
4.1.2.2 Ensure audit logs are not automatically deleted (Automated).....	298
4.1.2.3 Ensure system is disabled when audit logs are full (Automated)	299
4.1.3 Ensure changes to system administration scope (sudoers) is collected (Automated)	300
4.1.4 Ensure login and logout events are collected (Automated).....	302
4.1.5 Ensure session initiation information is collected (Automated)	304
4.1.6 Ensure events that modify date and time information are collected (Automated).....	306
4.1.7 Ensure events that modify the system's Mandatory Access Controls are collected (Automated)....	308
4.1.8 Ensure events that modify the system's network environment are collected (Automated).....	310
4.1.9 Ensure discretionary access control permission modification events are collected (Automated)	312
4.1.10 Ensure unsuccessful unauthorized file access attempts are collected (Automated).....	315
4.1.11 Ensure events that modify user/group information are collected (Automated)	318
4.1.12 Ensure successful file system mounts are collected (Automated).....	320
4.1.13 Ensure use of privileged commands is collected (Automated).....	322
4.1.14 Ensure file deletion events by users are collected (Automated).....	324
4.1.15 Ensure kernel module loading and unloading is collected (Automated).....	326
4.1.16 Ensure system administrator actions (sudolog) are collected (Automated)	328
4.1.17 Ensure the audit configuration is immutable (Automated)	330
4.2 Configure Logging.....	332
4.2.1 Configure rsyslog.....	333
4.2.1.1 Ensure rsyslog is installed (Automated)	334
4.2.1.2 Ensure rsyslog Service is enabled (Automated)	336
4.2.1.3 Ensure rsyslog default file permissions configured (Automated).....	338
4.2.1.4 Ensure logging is configured (Manual).....	340
4.2.1.5 Ensure rsyslog is configured to send logs to a remote log host (Automated).....	342
4.2.1.6 Ensure remote rsyslog messages are only accepted on designated log hosts. (Manual).....	344
4.2.2 Configure journald	346
4.2.2.1 Ensure journald is configured to send logs to rsyslog (Automated).....	347
4.2.2.2 Ensure journald is configured to compress large log files (Automated)	349
4.2.2.3 Ensure journald is configured to write logfiles to persistent disk (Automated).....	351
4.2.3 Ensure permissions on all logfiles are configured (Automated)	353
4.2.4 Ensure permissions on operating system resources (OSR) (Automated).....	355
4.3 Ensure logrotate is configured (Manual).....	357
4.4 Ensure logging of login success or failure (Automated)	358
5 Access, Authentication and Authorization	362
5.1 Configure cron.....	363
5.1.1 Ensure cron daemon is enabled (Automated).....	364
5.1.2 Ensure permissions on /etc/crontab are configured (Automated)	365
5.1.3 Ensure permissions on /etc/cron.hourly are configured (Automated).....	367
5.1.4 Ensure permissions on /etc/cron.daily are configured (Automated)	369
5.1.5 Ensure permissions on /etc/cron.weekly are configured (Automated).....	371
5.1.6 Ensure permissions on /etc/cron.monthly are configured (Automated)	373
5.1.7 Ensure permissions on /etc/cron.d are configured (Automated)	375
5.1.8 Ensure at/cron is restricted to authorized users (Automated)	377
5.2 SSH Server Configuration	379
5.2.1 Ensure permissions on /etc/ssh/sshd_config are configured (Automated)	380

5.2.2 Ensure SSH access is limited (Automated)	381
5.2.3 Ensure permissions on SSH private host key files are configured (Automated)	383
5.2.4 Ensure permissions on SSH public host key files are configured (Automated)	386
5.2.5 Ensure SSH LogLevel is appropriate (Automated)	389
5.2.6 Ensure SSH X11 forwarding is disabled (Automated)	391
5.2.7 Ensure SSH MaxAuthTries is set to 4 or less (Automated)	392
5.2.8 Ensure SSH IgnoreRhosts is enabled (Automated)	393
5.2.9 Ensure SSH HostbasedAuthentication is disabled (Automated)	394
5.2.10 Ensure SSH root login is disabled (Automated)	395
5.2.11 Ensure SSH PermitEmptyPasswords is disabled (Automated)	397
5.2.12 Ensure SSH PermitUserEnvironment is disabled (Automated)	398
5.2.13 Ensure SSH Idle Timeout Interval is configured (Automated)	399
5.2.14 Ensure SSH LoginGraceTime is set to one minute or less (Automated)	401
5.2.15 Ensure SSH warning banner is configured (Automated)	402
5.2.16 Ensure SSH PAM is enabled (Automated)	403
5.2.17 Ensure SSH AllowTcpForwarding is disabled (Automated)	404
5.2.18 Ensure SSH MaxStartups is configured (Automated)	406
5.2.19 Ensure SSH MaxSessions is set to 4 or less (Automated)	407
5.2.20 Ensure system-wide crypto policy is not over-ridden (Automated)	408
5.3 Configure authselect	409
5.3.1 Create custom authselect profile (Automated)	410
5.3.2 Select authselect profile (Automated)	411
5.3.3 Ensure authselect includes with-faillock (Automated)	412
5.4 Configure PAM	413
5.4.1 Ensure password creation requirements are configured (Automated)	414
5.4.2 Ensure lockout for failed password attempts is configured (Automated)	417
5.4.3 Ensure password reuse is limited (Automated)	420
5.4.4 Ensure password hashing algorithm is SHA-512 (Automated)	422
5.5 User Accounts and Environment	424
5.5.1 Set Shadow Password Suite Parameters	425
5.5.1.1 Ensure password expiration is 365 days or less (Automated)	426
5.5.1.2 Ensure minimum days between password changes is 7 or more (Automated)	428
5.5.1.3 Ensure password expiration warning days is 7 or more (Automated)	430
5.5.1.4 Ensure inactive password lock is 30 days or less (Automated)	432
5.5.1.5 Ensure all users last password change date is in the past (Automated)	434
5.5.2 Ensure system accounts are secured (Automated)	435
5.5.3 Ensure default user shell timeout is 900 seconds or less (Automated)	437
5.5.4 Ensure default group for the root account is GID 0 (Automated)	439
5.5.5 Ensure default user umask is 027 or more restrictive (Automated)	440
5.6 Ensure root login is restricted to system console (Manual)	442
5.7 Ensure access to the su command is restricted (Automated)	443
6 System Maintenance	445
6.1 System File Permissions	446
6.1.1 Audit system file permissions (Manual)	447
6.1.2 Ensure permissions on /etc/passwd are configured (Automated)	449
6.1.3 Ensure permissions on /etc/shadow are configured (Automated)	450
6.1.4 Ensure permissions on /etc/group are configured (Automated)	452
6.1.5 Ensure permissions on /etc/gshadow are configured (Automated)	453
6.1.6 Ensure permissions on /etc/passwd- are configured (Automated)	455
6.1.7 Ensure permissions on /etc/shadow- are configured (Automated)	456
6.1.8 Ensure permissions on /etc/group- are configured (Automated)	458

6.1.9 Ensure permissions on /etc/gshadow- are configured (Automated)	459
6.1.10 Ensure no world writable files exist (Automated)	461
6.1.11 Ensure no unowned files or directories exist (Automated)	463
6.1.12 Ensure no ungrouped files or directories exist (Automated)	465
6.1.13 Audit SUID executables (Manual)	467
6.1.14 Audit SGID executables (Manual)	469
6.1.15 Ensure udev rules have proper ownership (Manual)	471
6.2 User and Group Settings	473
6.2.1 Ensure password fields are not empty (Automated)	474
6.2.2 Ensure no legacy "+" entries exist in /etc/passwd (Automated)	475
6.2.3 Ensure root PATH Integrity (Automated)	476
6.2.4 Ensure no legacy "+" entries exist in /etc/shadow (Automated)	477
6.2.5 Ensure no legacy "+" entries exist in /etc/group (Automated)	478
6.2.6 Ensure root is the only UID 0 account (Automated)	479
6.2.7 Ensure users' home directories permissions are 750 or more restrictive (Automated)	480
6.2.8 Ensure users own their home directories (Automated)	483
6.2.9 Ensure users' dot files are not group or world writable (Automated)	485
6.2.10 Ensure no users have .forward files (Automated)	487
6.2.11 Ensure no users have .netrc files (Automated)	488
6.2.12 Ensure users' .netrc Files are not group or world accessible (Automated)	489
6.2.13 Ensure no users have .rhosts files (Automated)	492
6.2.14 Ensure all groups in /etc/passwd exist in /etc/group (Automated)	493
6.2.15 Ensure no duplicate UIDs exist (Automated)	494
6.2.16 Ensure no duplicate GIDs exist (Automated)	496
6.2.17 Ensure no duplicate user names exist (Automated)	498
6.2.18 Ensure no duplicate group names exist (Automated)	499
6.2.19 Ensure shadow group is empty (Automated)	500
6.2.20 Ensure all users' home directories exist (Automated)	501
6.3 Enable timeout support for login shells (Automated)	502
7 Hardware	505
7.1 Cryptography	506
7.1.1 Ensure correct assignment of cryptographic coprocessors and domains for EP11 tokens (Manual)	507
7.1.2 Ensure correct control domain assignment (Manual)	508
7.1.3 Ensure access to cryptographic devices is allowed only through device nodes created by zcryptctl (Manual)	509
7.1.4 Configure redundant APQNs and make sure they are on different adapters (Manual)	510
7.1.5 Enable CPACF feature (Manual)	511
7.2 I/O	513
7.2.1 Ensure cio_ignore is configured (Manual)	514
7.2.2 Enable SMART monitor (Manual)	515
7.2.3 Configure dm-crypt to use XTS-PAES with 256 bit keys (Manual)	516
7.3 Peripheral Component Interconnect (PCI) Devices	517
7.3.1 Ensure UID uniqueness checking is active (Manual)	518
7.3.2 NVMe setting for data resilience (Manual)	519
7.3.3 Ensure that unused PCI Devices are disabled (Manual)	521
7.4 Network	522
7.4.1 Configure SMC-D for network traffic isolation (Manual)	523
7.4.2 Ensure RoCE card redundancy settings are configured properly (Manual)	524
Appendix: Summary Table	525

<i>Appendix: Change History</i>	<i>544</i>

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 Red Hat Enterprise Linux derived Linux distribution systems running on s390 platforms.

The guidance within broadly assumes that operations are being performed as the root user. Operations performed using sudo instead of the root user 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.

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 Red Hat Enterprise Linux derived Linux distributions on s390 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.

- **Level 1 - Workstation**

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 workstations.

- **Level 2 - Workstation**

This profile extends the "Level 1 - Workstation" 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 workstations.

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

Rael Daruszka
Ron Colvin
Bill Erickson
Dave Billing
Dominic Pace
Elliot Anderson
Ely Pinto
Fredrik Silverskär
Joy Latten
Kirill Antonenko
Koen Laevens
Marcelo Cerri
Mark Birch
Martynas Brijunas
Michel Verbraak
Mike Thompson
Pradeep R B
Rakesh Jain
Robert Thomas
Tom Pietschmann
Vineetha Hari Pai
William E. Triest III
Anurag Pal
Bradley Hieber
Thomas Sjögren
James Trigg
Kenneth Karlsson
Mark Hesse
Pradeep Parameshwaran
Stephan Hartig
Robert BLESSING-HARTLEY
Horst Weber
Shaun Kelley
Eric Pinnell

Editor

Jonathan Lewis Christopherson
Eric Pinnell

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, make sure the data has been copied over to the new partition, unmount it and then remove the data from the directory that was in the old partition. Otherwise it will still consume space in the old partition that will be masked when the new filesystem is mounted. For example, if a system is in single-user mode with no filesystems mounted and the administrator adds a lot of data to the `/tmp` directory, this data will still consume space in `/` once the `/tmp` filesystem is mounted unless it is removed first.

1.1.1 Disable unused filesystems

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.

1.1.1.1 Ensure mounting of cramfs filesystems is disabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 server. If this filesystem type is not needed, disable it.

Audit:

Run the following commands and verify the output is as indicated:

```
# modprobe -n -v cramfs

install /bin/true
# lsmod | grep cramfs

<No output>
```

Remediation:

Edit or create a file in the `/etc/modprobe.d/` directory ending in `.conf`

Example: `vim /etc/modprobe.d/cramfs.conf`

and add the following line:

```
install cramfs /bin/true
```

Run the following command to unload the `cramfs` module:

```
# rmmod cramfs
```

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.			

1.1.1.2 Ensure mounting of vFAT filesystems is limited (Manual)

Profile Applicability:

- Level 2 - Workstation
- Level 2 - Server

Description:

The `vFAT` filesystem format is primarily used on older windows systems and portable USB drives or flash modules. It comes in three types `FAT12` , `FAT16` , and `FAT32` all of which are supported by the `vfat` kernel module.

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:

The `FAT` filesystem format is used by UEFI systems for the EFI boot partition. Disabling the `vfat` module can prevent boot on UEFI systems.

`FAT` filesystems are often used on portable USB sticks and other flash media which are commonly used to transfer files between workstations, removing VFAT support may prevent the ability to transfer files in this way.

Audit:

If utilizing UEFI the `vFAT` filesystem format is required. If this case, ensure that the `vFAT` filesystem is only used where appropriate

Run the following command

```
grep -E -i '\svfat\s' /etc/fstab
```

And review that any output is appropriate for your environment

If not utilizing UEFI

Run the following commands and verify the output is as indicated:

```
# modprobe -n -v vfat
```

```
install /bin/true
```

```
# lsmod | grep vfat
```

```
<No output>
```

Remediation:

Edit or create a file in the /etc/modprobe.d/ directory ending in .conf

Example: vim /etc/modprobe.d/vfat.conf

```
install vfat /bin/true
```

Run the following command to unload the `vfat` module:

```
# rmmod vfat
```

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.			

1.1.1.3 Ensure mounting of squashfs filesystems is disabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `squashfs` filesystem type is a compressed read-only Linux filesystem embedded in small footprint systems (similar to `cramfs`). 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.

Audit:

Run the following commands and verify the output is as indicated:

```
# modprobe -n -v squashfs

install /bin/true
# lsmod | grep squashfs

<No output>
```

Remediation:

Edit or create a file in the `/etc/modprobe.d/` directory ending in `.conf`

Example: `vim /etc/modprobe.d/squashfs.conf`

and add the following line:

```
install squashfs /bin/true
```

Run the following command to unload the `squashfs` module:

```
# rmmod squashfs
```

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	5.1 <u>Establish Secure Configurations</u> Maintain documented, standard security configuration standards for all authorized operating systems and software.	●	●	●

1.1.1.4 Ensure mounting of udf filesystems is disabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 commands and verify the output is as indicated:

```
# modprobe -n -v udf

install /bin/true
# lsmod | grep udf

<No output>
```

Remediation:

Edit or create a file in the `/etc/modprobe.d/` directory ending in `.conf`

Example: `vim /etc/modprobe.d/udf.conf`

and add the following line:

```
install udf /bin/true
```

Run the following command to unload the `udf` module:

```
# rmmod udf
```

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.			

1.1.2 Ensure /tmp is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 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 hardlink to a system `setuid` program and wait for it to be updated. Once the program was updated, the hardlink 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:

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

Running out of `/tmp` space is a problem regardless of what kind of filesystem lies under it, but in a default installation a disk-based `/tmp` will essentially have the whole disk available, as it 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.

`/tmp` utilizing tmpfs can be resized using the `size={size}` parameter on the Options line on the `tmp.mount` file

Audit:

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

```
# mount | grep -E '\s/tmp\s'

tmpfs on /tmp type tmpfs (rw,nosuid,nodev,noexec,relatime)
```

Run the following command and verify that tmpfs has been mounted to, or a system partition has been created for `/tmp`

```
# grep -E '\s/tmp\s' /etc/fstab | grep -E -v '^s*#'

tmpfs    /tmp     tmpfs    defaults,noexec,nosuid,nodev 0    0
```

OR

```
# systemctl is-enabled tmp.mount

enabled
```

Remediation:

Configure `/etc/fstab` as appropriate.

```
example:
tmpfs    /tmp     tmpfs    defaults,rw,nosuid,nodev,noexec,relatime 0 0
```

OR

Run the following commands to enable systemd `/tmp` mounting:

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

Edit `/etc/systemd/system/local-fs.target.wants/tmp.mount` to configure the `/tmp` mount:

```
[Mount]
What=tmpfs
Where=/tmp
Type=tmpfs
Options=mode=1777,strictatime,noexec,nodev,nosuid
```

References:

1. AJ Lewis, "LVM HOWTO", <http://tldp.org/HOWTO/LVM-HOWTO/>
2. <https://www.freedesktop.org/wiki/Software/systemd/APIFileSystems/>

Additional Information:

If an entry for /tmp exists in /etc/fstab it will take precedence over entries in the tmp.mount file

BUG 1667065* There is currently a bug in RHEL 8 when attempting to use systemd tmp.mount please reference link bellow

https://bugzilla.redhat.com/show_bug.cgi?id=1667065

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.			

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

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 attempt to create block or character special devices in `/tmp`.

Audit:

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

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

Remediation:

Edit the `/etc/fstab` file and add `nodev` to the fourth field (mounting options) for the `/tmp` partition. See the `fstab(5)` manual page for more information.

Run the following command to remount `/tmp`:

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

OR

Edit `/etc/systemd/system/local-fs.target.wants/tmp.mount` to add `nodev` to the `/tmp` mount options:

```
[Mount]
Options=mode=1777,strictatime,noexec,nodev,nosuid
```

Run the following command to remount `/tmp`:

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


CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>3.3 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	<u>14.6 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.			

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

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 if a `/tmp` partition exists
Run the following command and verify that nothing is returned:

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

Remediation:

Edit the `/etc/fstab` file and add `nosuid` to the fourth field (mounting options) for the `/tmp` partition. See the `fstab(5)` manual page for more information.

Run the following command to remount `/tmp`:

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

or

Edit `/etc/systemd/system/local-fs.target.wants/tmp.mount` to add `nosuid` to the `/tmp` mount options:

```
[Mount]
Options=mode=1777,strictatime,noexec,nodev,nosuid
```

Run the following command to remount `/tmp`:

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

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.			

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

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 if a `/tmp` partition exists
Run the following command and verify that nothing is returned:

```
# mount | grep -E '\s/tmp\s' | grep -v noexec
```

Remediation:

Edit the `/etc/fstab` file and add `noexec` to the fourth field (mounting options) for the `/tmp` partition. See the `fstab(5)` manual page for more information.

Run the following command to remount `/tmp` :

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

or

Edit `/etc/systemd/system/local-fs.target.wants/tmp.mount` to add `noexec` to the `/tmp` mount options:

```
[Mount]
Options=mode=1777,strictatime,noexec,nodev,nosuid
```

Run the following command to remount `/tmp` :

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

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.			

1.1.6 Ensure separate partition exists for /var (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

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:

Since the `/var` directory may contain world-writable files and directories, there is a risk of resource exhaustion if it is not bound to a separate partition.

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:

```
# mount | grep -E '\s/var\s'
/dev/xvdgl on /var type xfs (rw,relatime,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/>

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 multiuser 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.			
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.			

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

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

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

Rationale:

Since the `/var/tmp` directory is intended to be world-writable, there is a risk of resource exhaustion if it is not bound to a separate partition. In addition, making `/var/tmp` its own file system allows an administrator to set the `noexec` option on the mount, making `/var/tmp` useless for an attacker to install executable code. It would also prevent an attacker from establishing a hardlink to a system `setuid` program and wait for it to be updated. Once the program was updated, the hardlink 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:

```
# mount | grep /var/tmp  
<device> on /var/tmp type xfs (rw,nosuid,nodev,noexec,relatime)
```

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.

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.			

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

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 attempt to create block or character special devices in `/var/tmp`.

Audit:

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

```
# mount | grep -E '\s/var/tmp\s' | grep -v nodev
```

Remediation:

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

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

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.			

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

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 if a `/var/tmp` partition exists.
Run the following command and verify that nothing is returned:

```
# mount | grep -E '\s/var/tmp\s' | grep -v nosuid
```

Remediation:

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

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

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.			

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

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 if a `/var/tmp` partition exists.
Run the following command and verify that nothing is returned:

```
# mount | grep -E '\s/var/tmp\s' | grep -v noexec
```

Remediation:

Edit the `/etc/fstab` file and add `noexec` to the fourth field (mounting options) for the `/var/tmp` partition. See the `fstab(5)` manual page for more information.
Run the following command to remount `/var/tmp`:

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

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.			

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

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

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

Rationale:

There are two important reasons to ensure that system logs are stored on a separate partition: protection against resource exhaustion (since logs can grow quite large) and protection of audit data.

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:

```
# mount | grep /var/log  
  
/dev/xvdh1 on /var/log type xfs (rw,relatime,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/>

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.			

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

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

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

Rationale:

There are two important reasons to ensure that data gathered by `auditd` is stored on a separate partition: protection against resource exhaustion (since the `audit.log` file can grow quite large) and protection of audit data. The audit daemon calculates how much free space is left and performs actions based on the results. If other processes (such as `syslog`) consume space in the same partition as `auditd`, it may not perform as desired.

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:

```
# mount | grep /var/log/audit  
/dev/xvdi1 on /var/log/audit type xfs (rw,relatime,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/>

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 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.			

1.1.13 Ensure separate partition exists for /home (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

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

Rationale:

If the system is intended to support local users, create a separate partition for the `/home` directory to protect against resource exhaustion and restrict the type of files that can be stored under `/home`.

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:

```
# mount | grep /home  
  
/dev/xvdf1 on /home type xfs (rw,nodev,relatime,data=ordered)
```

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/>

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.			

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

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

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

Rationale:

Since the user partitions are not intended to support devices, set this option to ensure that users cannot attempt to create block or character special devices.

Audit:

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

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

Remediation:

Edit the `/etc/fstab` file and add `nodev` to the fourth field (mounting options) for the `/home` partition. See the `fstab(5)` manual page for more information.

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

Additional Information:

The actions in this recommendation refer to the `/home` partition, which is the default user partition that is defined in many distributions. If you have created other user partitions, it is recommended that the Remediation and Audit steps be applied to these partitions as well.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>3.3 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	<u>14.6 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.			

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

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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`:

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

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	<u>3.3 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	<u>14.6 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.			

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

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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`:

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

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	<u>3.3 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	<u>14.6 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.			

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

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 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 noexec
```

Remediation:

Edit the `/etc/fstab` file and add `noexec` 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`:

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

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	<u>3.3 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	<u>14.6 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.			

1.1.18 Ensure nodev option set on removable media partitions (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

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

Rationale:

Removable media containing character and block special devices could be used to circumvent security controls by allowing non-root users to access sensitive device files such as `/dev/kmem` or the raw disk partitions.

Audit:

Run the following command and verify that the `nodev` option is set on all removable media partitions.

```
# mount
```

Remediation:

Edit the `/etc/fstab` file and add `nodev` to the fourth field (mounting options) of all removable media partitions. Look for entries that have mount points that contain words such as `floppy` or `cdrom`. See the `fstab(5)` manual page for more information.

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.			

1.1.19 Ensure nosuid option set on removable media partitions (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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:

Run the following command and verify that the `nosuid` option is set on all removable media partitions.

```
# mount
```

Remediation:

Edit the `/etc/fstab` file and add `nosuid` to the fourth field (mounting options) of all removable media partitions. Look for entries that have mount points that contain words such as `floppy` or `cdrom`. See the `fstab(5)` manual page for more information.

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.			

1.1.20 Ensure noexec option set on removable media partitions (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 the removable media. This deters users from being able to introduce potentially malicious software on the system.

Audit:

Run the following command and verify that the `noexec` option is set on all removable media partitions.

```
# mount
```

Remediation:

Edit the `/etc/fstab` file and add `noexec` to the fourth field (mounting options) of all removable media partitions. Look for entries that have mount points that contain words such as `floppy` or `cdrom`. See the `fstab(5)` manual page for more information.

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.			

1.1.21 Ensure sticky bit is set on all world-writable directories (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

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:

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 command to verify no world writable directories exist without the sticky bit set:

```
# df --local -P | awk '{if (NR!=1) print $6}' | xargs -I '{}' find '{}' -xdev -type d \( -perm -0002 -a ! -perm -1000 \) 2>/dev/null
```

No output should be returned.

Remediation:

Run the following command to set the sticky bit on all world writable directories:

```
# df --local -P | awk '{if (NR!=1) print $6}' | xargs -I '{}' find '{}' -xdev -type d \( -perm -0002 -a ! -perm -1000 \) 2>/dev/null | xargs -I '{}' chmod a+t '{}'
```

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.			

1.1.22 Disable Automounting (Automated)

Profile Applicability:

- Level 1 - Server
- Level 2 - Workstation

Description:

`autofs` allows automatic mounting of devices, typically including CD/DVDs and USB drives.

Rationale:

With automounting enabled anyone with physical access could attach a USB drive or disc and have its contents available in system even if they lacked permissions to mount it themselves.

Impact:

The use of portable hard drives is very common for workstation users. If your organization allows the use of portable storage or media on workstations and physical access controls to workstations is considered adequate there is little value add in turning off automounting.

Audit:

Run the following command to verify `autofs` is not enabled:

```
# systemctl is-enabled autofs  
disabled
```

Verify result is not "enabled".

Remediation:

Run the following command to disable `autofs`:

```
# systemctl --now disable autofs
```


Additional Information:

Additional methods of disabling a service exist. Consult your distribution documentation for appropriate methods.

This control should align with the tolerance of the use of portable drives and optical media in the organization. On a server requiring an admin to manually mount media can be part of defense-in-depth to reduce the risk of unapproved software or information being introduced or proprietary software or information being exfiltrated. If admins commonly use flash drives and Server access has sufficient physical controls, requiring manual mounting may not increase security.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>10.3 Disable Autorun and Autoplay for Removable Media</u> Disable autorun and autoplay auto-execute functionality for removable media.			
v7	<u>8.4 Configure Anti-Malware Scanning of Removable Devices</u> Configure devices so that they automatically conduct an anti-malware scan of removable media when inserted or connected.			
v7	<u>8.5 Configure Devices Not To Auto-run Content</u> Configure devices to not auto-run content from removable media.			

1.1.23 Disable USB Storage (Automated)

Profile Applicability:

- Level 1 - Server
- Level 2 - Workstation

Description:

USB storage provides a means to transfer and store files insuring 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.

Audit:

Run the following commands and verify the output is as indicated:

```
# modprobe -n -v usb-storage  
  
install /bin/true  
# lsmod | grep usb-storage  
  
<No output>
```

Remediation:

Edit or create a file in the /etc/modprobe.d/ directory ending in .conf
Example: vim /etc/modprobe.d/usb-storage.conf
and add the following line:

```
install usb-storage /bin/true
```

Run the following command to unload the usb-storage module:

```
rmmod usb-storage
```

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	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	8.4 <u>Configure Anti-Malware Scanning of Removable Devices</u> Configure devices so that they automatically conduct an anti-malware scan of removable media when inserted or connected.			
v7	8.5 <u>Configure Devices Not To Auto-run Content</u> Configure devices to not auto-run content from removable media.			

1.1.24 Mount devices read-only that do not require write access (Optional) (Manual)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Devices that do not require write-access should be mounted to the filesystem with the read-only option enabled.

Rationale:

Inadvertent deletion or modification of devices can lead to data loss or corruption.

Audit:

Enter the command:

```
cat /etc/fstab
```

Verify that the `ro` option is given for all devices that do not require write-access.

Remediation:

Add the `ro` option in file

```
/etc/fstab
```

to all devices that do not require write-access.

Default Value:

```
rw
```

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>3.3 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	<u>14.6 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.			

1.2 Configure Software Updates

Fedora 28 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. Many large enterprises prefer to test patches on a non-production system before rolling out to production.

For the purpose of this benchmark, the requirement is to ensure that a patch management system is configured and maintained. The specifics on patch update procedures are left to the organization.

1.2.1 Ensure GPG keys are configured (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Most packages managers implement GPG key signing to verify package integrity during 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.

Audit:

Verify GPG keys are configured correctly for your package manager. Depending on the package management in use one of the following command groups may provide the needed information:

```
# rpm -q gpg-pubkey --qf '%{name}-%{version}-%{release} --> %{summary}\n'
```

Remediation:

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

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	7.3 <u>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	7.4 <u>Perform Automated Application Patch Management</u> Perform application updates on enterprise assets through automated patch management on a monthly, or more frequent, basis.			
v7	3.4 <u>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.			

1.2.2 Ensure *gpgcheck* is globally activated (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `gpgcheck` option, found in the main section of the `/etc/yum.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:

Run the following command and verify `gpgcheck` is set to '1':

```
# grep ^gpgcheck /etc/yum.conf  
  
gpgcheck=1
```

Run the following command and verify that all instances of `gpgcheck` returned are set to '1':

```
# grep ^gpgcheck /etc/yum/repos.d/*
```

Remediation:

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

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

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.			

1.2.3 Ensure package manager repositories are configured (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Systems need to have package manager repositories configured to ensure they 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:

```
# dnf repolist
```

Remediation:

Configure your package manager repositories according to site policy.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	7.3 <u>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	7.4 <u>Perform Automated Application Patch Management</u> Perform application updates on enterprise assets through automated patch management on a monthly, or more frequent, basis.			
v7	3.4 <u>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.			

1.3 Configure sudo

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.

sudo supports a plugin architecture for security policies and input/output logging. Third parties can develop and distribute their own policy and I/O logging plugins to work seamlessly with the sudo front end. The default security policy is sudoers, which is configured via the file `/etc/sudoers`.

1.3.1 Ensure sudo is installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 plugin architecture for security policies and input/output logging. Third parties can develop and distribute their own policy and I/O logging plugins to work seamlessly with the sudo front end. The default security policy is sudoers, which is configured via the file /etc/sudoers.

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:

```
# rpm -q sudo  
  
sudo-<VERSION>
```

Remediation:

Run the following command to install sudo

```
# dnf install sudo
```

References:

1. SUDO(8)

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
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.			

1.3.2 Ensure sudo commands use pty (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

sudo can be configured to run only from a psuedo-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:

editing the sudo configuration incorrectly can cause sudo to stop functioning.

Audit:

Verify that sudo can only run other commands from a psuedo-pty
Run the following command:

```
# grep -Ei '^s*Defaults\s+(\[^\#\]+,\s*)?use_pty' /etc/sudoers
/etc/sudoers.d/*

Defaults use_pty
```

Remediation:

edit the file `/etc/sudoers` or a file in `/etc/sudoers.d/` with `visudo -f <PATH TO FILE>` and add the following line:

```
Defaults use_pty
```

References:

1. SUDO(8)
2. VISUDO(8)

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 or 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
v7	5.1 <u>Establish Secure Configurations</u> Maintain documented, standard security configuration standards for all authorized operating systems and software.			

1.3.3 Ensure sudo log file exists (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

sudo can use a custom log file

Rationale:

A sudo log file simplifies auditing of sudo commands

Impact:

editing the sudo configuration incorrectly can cause sudo to stop functioning

Audit:

Verify that sudo has a custom log file configured

Run the following command:

```
# grep -Ei '^s*Defaults\s+([\#]+,\s*)?logfile=' /etc/sudoers
/etc/sudoers.d/*

Defaults logfile="/var/log/sudo.log"
```

Remediation:

edit the file `/etc/sudoers` or a file in `/etc/sudoers.d/` with `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)

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 <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.			

1.3.4 Ensure sudo-specific activity logging (Manual)

Profile Applicability:

- Level 1 - Server
- Level 2 - Server

Description:

You can apply the logging of sudo activity into OS log files. Another way is to accomplish sudo activity logging into a sudo-specific log file. Such a file must be retained as required by the base policy. If you use a sudo-specific log file, the file `/var/log/sudo.log` must exist.

Rationale:

The processing of privileged commands must be logged.

Audit:

Execute the `sudo -v` command as root user, if the **Sudo Configure** option contains one of the following entries:

```
--with-logging=both  
--with-logging=file
```

Then use the `ls` command to verify that the file

```
/var/log/sudo.log
```

exists.

Remediation:

For all Unix distributions you must use the IBM ESD-compiled sudo binary. For Linux, you must use the standard sudo binary which is included in your Linux distribution.

Default Value:

N/A

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.			

1.4 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.

1.4.1 Ensure AIDE is installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

AIDE takes a snapshot of filesystem state 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>

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).			●

1.4.2 Ensure filesystem integrity is regularly checked (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 to verify that aidecheck.service and aidecheck.timer are enabled and running

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

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

OR

Run the following commands to determine if there is a `cron` job scheduled to run the aide check.

```
# crontab -u root -l | grep aide
# grep -r aide /etc/cron.* /etc/crontab
```

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

Remediation:

Run the following commands:

```
# cp ./config/aidecheck.service /etc/systemd/system/aidecheck.service
# cp ./config/aidecheck.timer /etc/systemd/system/aidecheck.timer
# chmod 0644 /etc/systemd/system/aidecheck.*

# systemctl reenabale aidecheck.timer
# systemctl restart aidecheck.timer
# systemctl daemon-reload
```

OR

Run the following command:

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

Add the following line to the crontab:

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

References:

1. <https://github.com/konstruktoid/hardening/blob/master/config/aidecheck.service>
2. <https://github.com/konstruktoid/hardening/blob/master/config/aidecheck.timer>

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 <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).			●

1.5 Secure Boot Settings

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

1.5.1 Ensure a secure IPL configuration (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Ensure that your IPL configuration options in the `zipl.conf` configuration file do not disable the secure boot support.

Rationale:

Secure boot support provides additional security during the boot process by performing signature verification of boot components. Disabling this support would allow an attacker to undetectably inject malicious code into the kernel environment.

Audit:

Enter the following command:

```
grep secure /etc/zipl.conf
```

In the output, verify that the value for `secure` is not set to 0 as shown in the following example:

```
secure=auto
```

Remediation:

Set the value for `secure` in `/etc/zipl.conf` to `auto`.

Default Value:

```
auto
```

1.5.2 Ensure permissions on the boot loader configuration are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Ensure that the `zipl.conf` configuration files are only writable by the `root` user.

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.

Impact:

Not restricting write access to the `zipl.conf` configuration files allows also regular users to change boot loader configurations, including the disabling of security mechanisms and the booting of different kernel images.

Audit:

Enter the following commands:

```
# ls -l /etc/zipl.conf
# ls -l /boot/loader/entries
```

Make sure that files belong to the `root` user or group and that they are not world-writable.

Remediation:

To ensure the correct write permissions, enter the following commands:

```
# chown root:root /etc/zipl.conf
# chown root:root /boot/loader/entries/*
# chmod 600 /etc/zipl.conf
# chmod 600 /boot/loader/entries/*
```

Default Value:

All files belong to user or group `root` and are not world-writable.

Additional Information:

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

Replace `/boot/grub2/grub.cfg` and `/boot/grub2/grubenv` with the appropriate configuration file(s) for your environment

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v7	5.1 <u>Establish Secure Configurations</u> Maintain documented, standard security configuration standards for all authorized operating systems and software.			

1.5.3 Ensure boot loader password is set (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Setting the boot loader password will require that anyone rebooting the system must enter a password before being able to set command line boot parameters

Rationale:

Requiring a boot password upon execution of the boot loader will prevent an unauthorized user from entering boot parameters or changing the boot partition. This prevents users from weakening security (e.g. turning off SELinux at boot time).

Impact:

If password protection is enabled, only the designated superuser can edit a Grub 2 menu item by pressing "e" or access the GRUB 2 command line by pressing "c"

If GRUB 2 is set up to boot automatically to a password-protected menu entry the user has no option to back out of the password prompt to select another menu entry. Holding the SHIFT key will not display the menu in this case. The user must enter the correct username and password. If unable, the configuration files will have to be edited via the LiveCD or other means to fix the problem

You can add --unrestricted to the menu entries to allow the system to boot without entering a password. Password will still be required to edit menu items.

Audit:

Run the following command:

```
# grep "^s*GRUB2_PASSWORD" /boot/grub2/user.cfg  
GRUB2_PASSWORD=<encrypted-password>
```

Remediation:

Create an encrypted password with `grub2-setpassword` :

```
# grub2-setpassword
Enter password: <password>
Confirm password: <password>
```

Run the following command to update the `grub2` configuration:

```
# grub2-mkconfig -o /boot/grub2/grub.cfg
```

Additional Information:

This recommendation is designed around the `grub2` boot loader, if LILO or another boot loader is in use in your environment enact equivalent settings.

Replace `/boot/grub2/grub.cfg` with the appropriate grub configuration file for your environment

The superuser/user information and password do not have to be contained in the `/etc/grub.d/00_header` file. The information can be placed in any `/etc/grub.d` file as long as that file is incorporated into `grub.cfg`. The user may prefer to enter this data into a custom file, such as `/etc/grub.d/40_custom` so it is not overwritten should the Grub package be updated. If placing the information in a custom file, do not include the "cat << EOF" and "EOF" lines as the content is automatically added from these files.

1.5.4 Ensure authentication required for single user mode (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Single user mode (rescue mode) is used for recovery when the system detects an issue during boot or by manual selection from the boot loader.

Rationale:

Requiring authentication in single user mode (rescue mode) prevents an unauthorized user from rebooting the system into single user to gain root privileges without credentials.

Audit:

Run the following commands and verify that `/sbin/sulogin` or `/usr/sbin/sulogin` is used as shown:

```
# grep /systemd-sulogin-shell /usr/lib/systemd/system/rescue.service

ExecStart=-/usr/lib/systemd/systemd-sulogin-shell rescue
# grep /systemd-sulogin-shell /usr/lib/systemd/system/emergency.service

ExecStart=-/usr/lib/systemd/systemd-sulogin-shell emergency
```

Remediation:

Edit `/usr/lib/systemd/system/rescue.service` and add/modify the following line:

```
ExecStart=-/usr/lib/systemd/systemd-sulogin-shell rescue
```

Edit `/usr/lib/systemd/system/emergency.service` and add/modify the following line:

```
ExecStart=-/usr/lib/systemd/systemd-sulogin-shell emergency
```

1.5.5 Enable secure boot as IPL parameter (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Ensure that secure boot is enabled with the IPL parameters if supported by the system.

Rationale:

Secure boot is a security standard to help to make sure that a booting device uses only software that is trusted by the original equipment manufacturer (OEM).

Audit:

Read the sysfs attribute `/sys/firmware/ipl/secure` to check whether the Linux instance was IPLed with secure boot.

Issue the following command:

```
cat /sys/firmware/ipl/secure
```

```
1
```

If the value 1 is returned, Linux was IPLed with secure boot.

Remediation:

Run the `zipl` command with following options:

`-S` or `--secure`

SCSI IPL disk device for LPAR only: This command controls the format of the boot data that the `zipl` command writes to the IPL device. Takes the following values:

1. `auto` ==> Uses the secure-boot enabled format if the `zipl` command is issued on a mainframe with secure-boot support. This is the default.
2. `1` ==> Enforces the secure-boot enabled format regardless of mainframe support. Use this option to prepare boot devices for systems other than the one you are working on. Disks with this format can only be booted on z15 machines or later.
3. `0` ==> Enforces the traditional format that does not support secure boot, regardless of mainframe support. Disks with this format can be booted on all machines but cannot be used for secure boot.

Default Value:

```
1
```


1.6 Additional Process Hardening

1.6.1 Ensure core dumps are restricted (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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. The system provides the ability to set a soft limit for core dumps, but this can be overridden by the user.

Rationale:

Setting a hard limit on core dumps prevents users from overriding the soft variable. If core dumps are required, consider setting limits for user groups (see `limits.conf(5)`). In addition, setting the `fs.suid_dumpable` variable to 0 will prevent setuid programs from dumping core.

Audit:

Run the following commands and verify output matches:

```
# grep -E "^s*\s*\s+hard\s+core" /etc/security/limits.conf
/etc/security/limits.d/*

* hard core 0
# sysctl fs.suid_dumpable

fs.suid_dumpable = 0
# grep "fs\.suid_dumpable" /etc/sysctl.conf /etc/sysctl.d/*

fs.suid_dumpable = 0
```

Run the following command to check if systemd-coredump is installed:

```
# systemctl is-enabled coredump.service
```

if `enabled` or `disabled` is returned `systemd-coredump` is installed

Remediation:

Add the following line to `/etc/security/limits.conf` or a `/etc/security/limits.d/*` file:

```
* hard core 0
```

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

```
fs.suid_dumpable = 0
```

Run the following command to set the active kernel parameter:

```
# sysctl -w fs.suid_dumpable=0
```

If `systemd-coredump` is installed:

edit `/etc/systemd/coredump.conf` and add/modify the following lines:

```
Storage=none  
ProcessSizeMax=0
```

Run the command:

```
systemctl daemon-reload
```

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

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 commands and verify output matches:

```
# sysctl kernel.randomize_va_space

kernel.randomize_va_space = 2
# grep "kernel\.randomize_va_space" /etc/sysctl.conf /etc/sysctl.d/*

kernel.randomize_va_space = 2
```

Remediation:

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

```
kernel.randomize_va_space = 2
```

Run the following command to set the active kernel parameter:

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

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	10.5 <u>Enable Anti-Exploitation Features</u> 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 <u>Enable Operating System Anti-Exploitation Features/ Deploy Anti-Exploit Technologies</u> 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.			

1.7 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.

1.7.1 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 28 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 28 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>

1.7.1.1 Ensure SELinux is installed (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

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
```

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.			

1.7.1.2 Ensure SELinux is not disabled in boot loader configuration (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

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.

Audit:

Run the following command and verify that no linux line has the `selinux=0` or `enforcing=0` parameters set:

```
# grep -E 'kernelopts=(\S+\s+)*(selinux=0|enforcing=0)+\b' /boot/grub2/grubenv
```

Nothing should be returned

Remediation:

Edit `/etc/default/grub` and remove all instances of `selinux=0` and `enforcing=0` from all `CMDLINE_LINUX` parameters:

```
GRUB_CMDLINE_LINUX_DEFAULT="quiet"

GRUB_CMDLINE_LINUX=""
```

Run the following command to update the `grub2` configuration:

```
# grub2-mkconfig -o /boot/grub2/grub.cfg
```

Additional Information:

This recommendation is designed around the grub 2 boot loader, if LILO or another boot loader is in use in your environment enact equivalent settings.

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.			

1.7.1.3 Ensure SELinux policy is configured (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

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
```

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.			

1.7.1.4 Ensure the SELinux state is set to 'enforcing' (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

The SELinux state should be set to `enforcing` at system boot time by the boot loader. During the boot process, the boot loader is responsible for starting the execution of the kernel and passing options to it. The default Red Hat Enterprise Linux 8 boot loader for IBM Z systems is called zipl.

Rationale:

SELinux must be enabled at boot time to ensure that the controls it provides are always in effect.

Impact:

Setting the SELinux state to `enforcing` ensures that SELinux can confine potentially compromised processes that might perform actions opposite to the security policy. The `enforcing` option is designed to prevent compromised processes from causing damage to the system, or to prevent that their privileges are even further elevated.

Audit:

Run the following commands and ensure output matches:

```
# grep -E '^s*SELINUX=enforcing' /etc/selinux/config
SELINUX=enforcing

# sestatus
SELinux status: enabled
Current mode: enforcing
Mode from config file: enforcing
```

Ensure that the `getsestatus` command returns `enforced`.

Remediation:

In the file

```
/etc/selinux/config
```

add or correct the SELINUX parameter to configure the system to boot into enforcing mode:

```
SELINUX=enforcing
```

Default Value:

```
SELINUX=enforcing
```

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.			

1.7.1.5 Ensure no unconfined services exist (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

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 not output is produced:

```
# ps -eZ | grep unconfined_service_t
```

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.

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	<u>3.3 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	<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.			

1.7.1.6 Ensure SETroubleshoot is not installed (Automated)

Profile Applicability:

- Level 2 - 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
```

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.			

1.7.1.7 Ensure the MCS Translation Service (mcstrans) is not installed (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

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
```

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.			

1.7.1.8 Ensure SELinux is not disabled in zipl boot loader (Manual)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

To ensure that SELinux is not disabled at boot time, check that no boot entry in

```
/boot/loader/entries/*.conf
```

has `selinux=0` included in its options.

Make sure that

```
/etc/zipl.conf
```

does not contain `image = setting`, because Red Hat Enterprise Linux 8 adheres to the **Boot Loader Specification (BLS)**.

Additionally, run the `zipl` command so that

```
/boot/bootmap
```

is updated.

Rationale:

Disabling a major host protection feature - such as SELinux - at boot time, prevents this feature from confining unwanted system services right from the beginning. Furthermore, it increases the probability that the protection feature will permanently remain disabled during system operation.

Audit:

Ensure the following:

```
/boot/loader/entries/*.conf
```

includes the setting: `selinux=0`

Also check the following:

```
/etc/zipl.conf
```

does not contain entry: `image = setting`.

Remediation:

Run the `zipl` command to update

```
/boot/bootmap
```

Default Value:

```
selinux=0
```

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.			

1.8 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/>

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.

1.8.1 Command Line Warning Banners

The `/etc/motd`, `/etc/issue`, and `/etc/issue.net` files govern warning banners for standard command line logins for both local and remote users.

1.8.1.1 Ensure message of the day is configured properly (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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
```


1.8.1.2 Ensure local login warning banner is configured properly (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 users only. All activity may be monitored and reported." > /etc/issue
```

1.8.1.3 Ensure remote login warning banner is configured properly (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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
```

1.8.1.4 Ensure permissions on /etc/motd are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 /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
```

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.			

1.8.1.5 Ensure permissions on /etc/issue are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 /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
```

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.			

1.8.1.6 Ensure permissions on /etc/issue.net are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 /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
```

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.			

1.8.2 Ensure GDM login banner is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

GDM is the GNOME Display Manager which handles graphical login for GNOME based systems.

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.

Audit:

If GDM is installed on the system verify that `/etc/gdm3/greeter.dconf-defaults` file exists and contains the following:

```
[org/gnome/login-screen]
banner-message-enable=true
banner-message-text='<banner message>'
```

Remediation:

Edit or create the file `/etc/gdm3/greeter.dconf-defaults` and add the following:

```
[org/gnome/login-screen]
banner-message-enable=true
banner-message-text='Authorized uses only. All activity may be monitored and
reported.'
```

Additional Information:

Additional options and sections may appear in the `/etc/dconf/db/gdm.d/01-banner-message` file.

If a different GUI login service is in use, consult your documentation and apply an equivalent banner.

1.9 Ensure updates, patches, and additional security software are installed (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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
```

Remediation:

Use your package manager to update all packages on the system according to site policy.

The following command will install all available updates:

```
# dnf update
```

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.			

1.10 Ensure system-wide crypto policy is not legacy (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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:

The default system-wide cryptographic policy in Oracle Linux 8 does not allow communication using older, insecure protocols. For environments that require to be compatible with Oracle Linux 5 and in some cases also with earlier releases, the less secure LEGACY policy level is available.

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>

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.			

1.11 Ensure system-wide crypto policy is *FUTURE* or *FIPS* (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

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

FUTURE: Is a conservative security level that is believed to withstand any near-term future attacks. This level does not allow the use of SHA-1 in signature algorithms. The RSA and Diffie-Hellman parameters are accepted if larger than 3071 bits. The level provides at least 128-bit security

FIPS: Conforms to the FIPS 140-2 requirements. This policy is used internally by the fips-mode-setup(8) tool which can switch the system into the FIPS 140-2 compliance mode. The level provides at least 112-bit security

Impact:

The system-wide cryptographic policy in Red Hat Enterprise Linux 8 does not allow communication using older, insecure protocols. For environments that require to be compatible with Oracle Linux 5 and in some cases also with earlier releases, the less secure LEGACY policy level is available.

Audit:

Run the following command to verify that the system-wide crypto policy is Future or FIPS

```
# grep -E -i '^s*(FUTURE|FIPS)\s*(\s+#+.*)?$' /etc/crypto-policies/config
```

Verify that either `FUTURE` or `FIPS` is returned

Remediation:

Run the following command to change the system-wide crypto policy

```
# update-crypto-policies --set FUTURE
```

OR

To switch the system to FIPS mode, run the following command:

```
# fips-mode-setup --enable
```

Default Value:

DEFAULT

References:

1. CRYPTO-POLICIES(7)
2. <https://access.redhat.com/articles/3642912#what-polices-are-provided-1>

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.			

1.12 Ensure that the kernel has loaded all s390-specific crypto modules (Manual)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

The s390-specific kernel modules are versions of the generic cryptographic modules which are optimized for security and performance on the IBM Z platform. It is recommended that they are loaded.

Rationale:

Generic cryptographic modules may not exploit s390-specific hardware features and lead to negative impact for security and performance.

Audit:

Enter the following command and make sure all listed modules are loaded:

```
$ lsmod | grep s390
s390_trng                16384  0
crc32_vx_s390            16384  4
ghash_s390               16384  0
aes_s390                 24576  0
des_s390                 20480  0
libdes                   28672  1 des_s390
sha3_512_s390            16384  0
sha3_256_s390            16384  0
sha512_s390              16384  0
sha256_s390              16384  0
sha1_s390                16384  0
sha_common               16384  5 sha3_256_s390,sha512_s390,
                        sha256_s390,sha1_s390,sha3_512_s390
```

Remediation:

Install and IPL a kernel which has all required modules enabled.

Default Value:

Modules are loaded by default.

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 inetd Services

inetd is a super-server daemon that provides internet services and passes connections to configured services. While not commonly used inetd and any unneeded inetd based services should be disabled if possible.

2.1.1 Ensure xinetd is not installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The eXtended InterNET Daemon (`xinetd`) is an open source super daemon that replaced the original `inetd` daemon. The `xinetd` daemon listens for well known services and dispatches the appropriate daemon to properly respond to service requests.

Rationale:

If there are no `xinetd` services required, it is recommended that the package be removed.

Audit:

Run the following command to verify `xinetd` is not installed:

```
# rpm -q xinetd  
  
package xinetd is not installed
```

Remediation:

Run the following command to remove `xinetd`:

```
# dnf remove xinetd
```

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.		●	●

2.2 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 they be disabled or deleted from the system to reduce the potential attack surface.

2.2.1 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 chrony.

2.2.1.1 Ensure time synchronization is in use (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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.

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:

On physical systems or virtual systems where host based time synchronization is not available verify that chrony is installed.

Run the following command to verify that chrony is installed

```
# rpm -q chrony  
  
chrony-<VERSION>
```

On virtual systems where host based time synchronization is available consult your virtualization software documentation and verify that host based synchronization is in use.

Remediation:

On physical systems or virtual systems where host based time synchronization is not available install `chrony`:

Run the following command to install chrony:

```
# dnf install chrony
```

On virtual systems where host based time synchronization is available consult your virtualization software documentation and setup host based synchronization.

Additional Information:

`systemd-timesyncd` is part of `systemd`. Some versions of `systemd` have been compiled without `systemd-timesyncd`. On these distributions, `chrony` or `NTP` should be used instead of `systemd-timesyncd`.

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.			

2.2.1.2 Ensure chrony is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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.

This recommendation only applies if `chrony` is in use on the system.

Audit:

Run the following command and verify remote server is configured properly:

```
# grep -E "^(server|pool)" /etc/chrony.conf  
  
server <remote-server>
```

Multiple servers may be configured.

Run the following command and verify the first field for the `chronyd` process is `chrony`:

```
# ps -ef | grep chronyd  
  
chrony      491      1  0 20:32 ?          00:00:00 /usr/sbin/chronyd
```

Remediation:

Add or edit server or pool lines to `/etc/chrony.conf` as appropriate:

```
server <remote-server>
```

Configure `chrony` to run as the `chrony` user

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.			

2.2.2 Ensure X Window System is not installed (Automated)

Profile Applicability:

- Level 1 - 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, if provided by your distribution.

Audit:

Run the following command to Verify X Windows System is not installed.

```
# rpm -qa xorg-x11*
```

Remediation:

Run the following command to remove the X Windows System packages.

```
# dnf remove xorg-x11*
```


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.			

2.2.3 Ensure rsync service is not enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `rsyncd` service can be used to synchronize files between systems over network links.

Rationale:

The `rsyncd` service presents a security risk as it uses unencrypted protocols for communication.

Audit:

Run the following command to verify `rsyncd` is not enabled:

```
# systemctl is-enabled rsyncd  
  
disabled
```

Verify result is not "enabled"

Remediation:

Run the following command to disable `rsyncd`:

```
# systemctl --now disable rsyncd
```

Additional Information:

Additional methods of disabling a service exist. Consult your distribution documentation for appropriate methods.

On some distributions the `rsync` service is known as `rsync`, not `rsyncd`.

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.			

2.2.4 Ensure Avahi Server is not enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 disable the service to reduce the potential attack surface.

Audit:

Run the following command to verify the `avahi-daemon` is not enabled:

```
# systemctl is-enabled avahi-daemon  
disabled
```

Verify result is not "enabled".

Remediation:

Run the following command to disable `avahi-daemon`:

```
# systemctl --now disable avahi-daemon
```

Additional Information:

Additional methods of disabling a service exist. Consult your distribution documentation for appropriate methods.

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.			

2.2.5 Ensure SNMP Server is not enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

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 SNMP v1, which transmits data in the clear and does not require authentication to execute commands. Unless absolutely necessary, it is recommended that the SNMP service not be used. If SNMP is required the server should be configured to disallow SNMP v1.

Audit:

Run the following command to verify `snmpd` is not enabled:

```
# systemctl is-enabled snmpd  
disabled
```

Verify result is not "enabled"

Remediation:

Run the following command to disable `snmpd`:

```
# systemctl --now disable snmpd
```

Additional Information:

Additional methods of disabling a service exist. Consult your distribution documentation for appropriate methods.

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.			

2.2.6 Disable the pcnfsd service (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

The `pcnfsd` service is an authentication and printing program, which uses NFS to provide file transfer services. If this service is not required to support certain applications, it must be denied by the system settings for service prevention, or this `pcnfsd` service must be completely removed from the system.

Rationale:

The `pcnfsd` service is vulnerable and exploitable and permits attackers to compromise the machine both locally and remotely. If PC NFS clients are nonetheless required within the environment, `Samba` is recommended as an alternative software solution. The `pcnfsd` daemon predates Microsoft's release of SMB (`Samba`) specifications. Standard Linux distributions do not provide this package, but it is readily available under third party sites.

Audit:

Run the following command to verify if `pcnfsd` is installed:

```
rpm -qa pcnfsd
```

This command should return nothing.
If it returns something like

```
pcnfsd-xxx
```

it must be removed.
(xxx stands for the version information,
which may vary.)

Remediation:

Run the following command to remove the `pcnfsd` daemon from the system:

```
rpm -e pcnfsd
```

Default Value:

N/A

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.			

2.2.7 Disable the tcpspray service (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

The `sprayd` service is used as a tool to generate UDP packets for testing and diagnosing network problems.

If this service is not required to support certain applications, it must be denied by the system settings for service prevention, or this `sprayd` service must be completely removed from the system.

Rationale:

The `tcpspray` service must be disabled if you are not running NFS, as it can be used by attackers in a Distributed Denial of Service (DDoS) attack. Standard Linux distributions do not provide this package, but it is available under third party sites.

Audit:

Run the following command to verify if `tcpspray` is installed:

```
rpm -qa tcpspray
```

This command should return nothing. If it returns something like

```
tcpspray-1.1-0.a.2.el7.rf.x86_64
```

it must be removed. The version information may vary.

Remediation:

Run the following command to remove the `tcpspray` service from the system:

```
rpm -e tcpspray
```

Default Value:

N/A

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.			

2.2.8 Ensure HTTP Proxy Server is not enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Squid is a standard proxy server used in many distributions and environments.

Rationale:

If there is no need for a proxy server, it is recommended that the squid proxy be deleted to reduce the potential attack surface.

Audit:

Run the following command to verify `squid` is not enabled:

```
# systemctl is-enabled squid  
  
disabled
```

Verify result is not "enabled"

Remediation:

Run the following command to disable `squid`:

```
# systemctl --now disable squid
```

Additional Information:

Additional methods of disabling a service exist. Consult your distribution documentation for appropriate methods.

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.			

2.2.9 Ensure Samba is not enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

With the Samba daemon, system administrators can configure their Linux systems to share file systems and directories with Windows desktops. Samba advertises the file systems and directories using the Server Message Block (SMB) protocol. Windows desktop users can 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, it is recommended to disable this service to reduce the potential attack surface. If this service is required on some systems, document details in agreed to value.

Audit:

Run the following command to verify that the Samba daemon is not enabled:

```
# systemctl is-enabled smb  
  
disabled
```

Verify that the result is `disabled`.

Remediation:

Run the following command to disable the Samba daemon:

```
# systemctl --now disable smb
```

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.			

2.2.10 Ensure IMAP and POP3 server is not enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`dovecot` is an open source IMAP and POP3 server for Linux based systems, which must be disabled.

Rationale:

Unless POP3 or IMAP servers, or both, are necessarily to be provided by this system, it is recommended that they are disabled to reduce the potential attack surface.

Audit:

Run the following command to verify `dovecot` servers are not enabled:

```
# systemctl is-enabled dovecot  
  
disabled
```

Verify that the result is `disabled`.

Remediation:

Run the following command to disable `dovecot` IMAP and POP3 servers:

```
# systemctl disable dovecot
```

Additional Information:

Additional methods of disabling a service exist. Consult your distribution documentation for appropriate methods.

Several IMAP/POP3 servers exist and can use other service names. `courier-imap` and `cyrus-imap` are example services that provide a mail server. These and other services should also be audited.

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.		●	●

2.2.11 Ensure HTTP server is not enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

HTTP or 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 package be deleted to reduce the potential attack surface.

Audit:

Run the following command to verify `httpd` is not enabled:

```
# systemctl is-enabled httpd  
  
disabled
```

Verify result is not "enabled".

Remediation:

Run the following command to disable `httpd`:

```
# systemctl --now disable httpd
```

Additional Information:

Additional methods of disabling a service exist. Consult your distribution documentation for appropriate methods.

Several `httpd` servers exist and can use other service names. `apache`, `apache2`, `lighttpd`, and `nginx` are example services that provide an HTTP server. These and other services should also be audited.

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.			

2.2.12 Ensure FTP Server is not enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The File Transfer Protocol (FTP) provides networked computers with the ability to transfer files.

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 deleted to reduce the potential attack surface.

Audit:

Run the following command to verify `vsftpd` is not enabled:

```
# systemctl is-enabled vsftpd  
  
disabled
```

Verify result is not "enabled".

Remediation:

Run the following command to disable `vsftpd`:

```
# systemctl --now disable vsftpd
```

Additional Information:

Additional methods of disabling a service exist. Consult your distribution documentation for appropriate methods.

Additional FTP servers also exist and should be audited.

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.			

2.2.13 Ensure DNS Server is not enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 deleted to reduce the potential attack surface.

Audit:

Run the following command to verify `named` is not enabled:

```
# systemctl is-enabled named  
  
disabled
```

Verify result is not "enabled".

Remediation:

Run the following command to disable `named` :

```
# systemctl --now disable named
```

Additional Information:

Additional methods of disabling a service exist. Consult your distribution documentation for appropriate methods.

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.			

2.2.14 Ensure NFS is not enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 export NFS shares, it is recommended that the NFS be disabled to reduce the remote attack surface.

Audit:

Run the following command to verify `nfs` is not enabled:

```
# systemctl is-enabled nfs  
  
disabled
```

Verify result is not "enabled".

Remediation:

Run the following commands to disable `NFS`:

```
# systemctl --now disable nfs
```

Additional Information:

Additional methods of disabling a service exist. Consult your distribution documentation for appropriate methods.

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.			

2.2.15 Ensure RPC is not enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The rpcbind service maps Remote Procedure Call (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.

Rationale:

If the system does not require rpc based services, it is recommended that rpcbind be disabled to reduce the remote attack surface.

Impact:

Because RPC-based services rely on rpcbind to make all connections with incoming client requests, rpcbind must be available before any of these services start

Audit:

Run the following command to verify rpcbind is not enabled:

```
# systemctl is-enabled rpcbind  
disabled
```

Verify result is not "enabled".

Remediation:

Run the following commands to disable rpcbind :

```
# systemctl --now disable rpcbind
```

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.			

2.2.16 Ensure LDAP server is not enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 server, it is recommended that the software be disabled to reduce the potential attack surface.

Audit:

Run the following commands to verify `slapd` is not enabled:

```
# systemctl is-enabled slapd  
  
disabled
```

Verify result is not "enabled".

Remediation:

Run the following command to disable `slapd`:

```
# systemctl --now disable slapd
```

References:

1. For more detailed documentation on OpenLDAP, go to the project homepage at <http://www.openldap.org>.

Additional Information:

Additional methods of disabling a service exist. Consult your distribution documentation for appropriate methods.

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.			

2.2.17 Ensure DHCP Server is not enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 this service be deleted to reduce the potential attack surface.

Audit:

Run the following command to verify `dhcpcd` is not enabled:

```
# systemctl is-enabled dhcpcd  
  
disabled
```

Verify result is not "enabled".

Remediation:

Run the following command to disable `dhcpcd`:

```
# systemctl --now disable dhcpcd
```

References:

1. More detailed documentation on DHCP is available at <http://www.isc.org/software/dhcp>.

Additional Information:

Additional methods of disabling a service exist. Consult your distribution documentation for appropriate methods.

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.			

2.2.18 Ensure CUPS is not enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 2 - Workstation

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 disabled to reduce the potential attack surface.

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 enabled:

```
# systemctl is-enabled cups  
disabled
```

Verify result is not "enabled".

Remediation:

Run the following command to disable `cups`:

```
# systemctl --now disable cups
```

References:

1. More detailed documentation on CUPS is available at the project homepage at <http://www.cups.org>.

Additional Information:

Additional methods of disabling a service exist. Consult your distribution documentation for appropriate methods.

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.			

2.2.19 Ensure NIS Server is not enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The Network Information Service (NIS) (formally known as Yellow Pages) is a client-server directory service protocol for distributing system configuration files. The NIS server is a collection of programs that allow for the distribution of configuration files.

Rationale:

The NIS service is inherently an insecure system that has been vulnerable to DOS attacks, buffer overflows and has poor authentication for querying NIS maps. NIS generally has been replaced by such protocols as Lightweight Directory Access Protocol (LDAP). It is recommended that the service be disabled and other, more secure services be used.

Audit:

Run the following command to verify `ypserv` is not enabled:

```
# systemctl is-enabled ypserv  
disabled
```

Verify result is not "enabled"

Remediation:

Run the following command to disable `ypserv`:

```
# systemctl --now disable ypserv
```

Additional Information:

Additional methods of disabling a service exist. Consult your distribution documentation for appropriate methods.

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.			

2.2.20 Ensure mail transfer agent is configured for local-only mode (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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.

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 -E ':25\s' | grep -E -v '\s(127.0.0.1|::1):25\s'
```

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
```

Additional Information:

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.

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.			

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 NIS Client is not installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The Network Information Service (NIS), formerly known as Yellow Pages, is a client-server directory service protocol used to distribute system configuration files. The NIS client (`ypbind`) was used to bind a machine to an NIS server and receive the distributed configuration files.

Rationale:

The NIS service is inherently an insecure system that has been vulnerable to DOS attacks, buffer overflows and has poor authentication for querying NIS maps. NIS generally has been replaced by such protocols as Lightweight Directory Access Protocol (LDAP). It is recommended that the service be removed.

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:

Verify `ypbind` is not installed.

Run the following command:

```
# rpm -q ypbind  
  
package ypbind is not installed
```

Remediation:

Run the following command to Uninstall `ybind`

```
# dnf remove ybind
```

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	●	●	●

2.3.2 Ensure telnet client is not installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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:

Verify `telnet` is not installed.

Run the following command:

```
# rpm -q telnet  
package telnet is not installed
```

Remediation:

Run the following command to uninstall `telnet`

```
# dnf remove telnet
```

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>4.5 Use Multifactor Authentication For All Administrative Access</u> Use multi-factor authentication and encrypted channels for all administrative account access.			

2.3.3 Ensure LDAP client is not installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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:

Verify openldap-clients is not installed.

Run the following command:

```
# rpm -q openldap-clients  
package openldap-clients is not installed
```

Remediation:

Run the following command to uninstall openldap-clients.

```
# dnf remove openldap-clients
```

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.			

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.

3.1 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.

3.1.1 Ensure IP forwarding is disabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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. Setting these flags to 0 ensures that a system is not allowed to forward packets.

You can specify this TCP/IP stack setting either in the file `/etc/sysctl.conf`, or in some other control file in directory `/etc/sysctl.d`. Or you can just use the default (`ip_forward = 0`) to not allow forwarding of packets.

Rationale:

Setting the flags 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 command and verify output matches:

```
# sysctl net.ipv4.ip_forward

net.ipv4.ip_forward = 0
# grep -E -s "^s*net\.ipv4\.ip_forward\s*=\s*1" /etc/sysctl.conf
/etc/sysctl.d/*.conf /usr/lib/sysctl.d/*.conf /run/sysctl.d/*.conf

No value should be returned
# sysctl net.ipv6.conf.all.forwarding

net.ipv6.conf.all.forwarding = 0
# grep -E -s "^s*net\.ipv6\.conf\.all\.forwarding\s*=\s*1" /etc/sysctl.conf
/etc/sysctl.d/*.conf /usr/lib/sysctl.d/*.conf /run/sysctl.d/*.conf

No value should be returned
```

Remediation:

Run the following commands to restore the default parameters and set the active kernel parameters:

```
# grep -Els "^s*net\.ipv4\.ip_forward\s*=\s*1" /etc/sysctl.conf
/etc/sysctl.d/*.conf /usr/lib/sysctl.d/*.conf /run/sysctl.d/*.conf | while
read filename; do sed -ri "s/^s*(net\.ipv4\.ip_forward\s*) (=) (\s*\S+\b).*$/#
*REMOVED* \1/" $filename; done; sysctl -w net.ipv4.ip_forward=0; sysctl -w
net.ipv4.route.flush=1
# grep -Els "^s*net\.ipv6\.conf\.all\.forwarding\s*=\s*1" /etc/sysctl.conf
/etc/sysctl.d/*.conf /usr/lib/sysctl.d/*.conf /run/sysctl.d/*.conf | while
read filename; do sed -ri
"s/^s*(net\.ipv6\.conf\.all\.forwarding\s*) (=) (\s*\S+\b).*$/# *REMOVED* \1/"
$filename; done; sysctl -w net.ipv6.conf.all.forwarding=0; sysctl -w
net.ipv6.route.flush=1
```


3.1.2 Ensure packet redirect sending is disabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 commands and verify output matches:

```
# sysctl net.ipv4.conf.all.send_redirects

net.ipv4.conf.all.send_redirects = 0
# sysctl net.ipv4.conf.default.send_redirects

net.ipv4.conf.default.send_redirects = 0
# grep "net\.ipv4\.conf\.all\.send_redirects" /etc/sysctl.conf
/etc/sysctl.d/*

net.ipv4.conf.all.send_redirects = 0
# grep "net\.ipv4\.conf\.default\.send_redirects" /etc/sysctl.conf
/etc/sysctl.d/*

net.ipv4.conf.default.send_redirects= 0
```

Remediation:

Set the following parameters in `/etc/sysctl.conf` or a `/etc/sysctl.d/*` file:

```
net.ipv4.conf.all.send_redirects = 0
net.ipv4.conf.default.send_redirects = 0
```

Run the following commands 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
```

3.1.3 Enable SMC-D pre-load libraries (Manual)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

If SMC-D is used for network communications between LPARs, you should load the SMC-D pre-load libraries by default on Linux on Z.

Rationale:

The pre-load libraries provide all necessary security defaults for a HiperSockets connection.

Audit:

Check if the environment variable is set to

```
LD_PRELOAD=libsmc-preload.so
```

In addition, you can run the

```
smc_run
```

command to see the the pre-load library status.

Remediation:

For binary applications, use the pre-load library `libsmc-preload.so` from the `smc-tools` package.

To do so, either run your application using the command `smc_run` from the same package, or set the environment variable to

```
LD_PRELOAD=libsmc-preload.so
```

respectively.

Default Value:

```
LD_PRELOAD=libsmc-preload.so
```

3.1.4 Enable VEPA mode traffic for OSA network devices (Manual)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

It is recommended to enable VEPA mode in order to redirect traffic between the same OSA devices to the switch.

Rationale:

Only if the device is set in VEPA mode, the port will offload switching functionality to the external entity as described in 802.1Qbg.

Audit:

Enter the following command:

```
cat /sys/devices/qeth/<devno>/isolation
```

Verify that the value for this attribute is set to `forward`.

Remediation:

Enter the following command:

```
echo forward >/sys/devices/qeth/<devno>/isolation
```

Note: VEPA mode requires hardware support from the switch.

Above command will fail if no support is available.

Default Value:

`none` for OSA devices.

3.1.5 Enable traffic isolation between OSA cards (Manual)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

When using networking via OSA devices, it is recommended that for each OSA device the traffic isolation is enabled.

Rationale:

If the isolation parameter of OSA devices is not set to **drop**, it is possible for other OSA devices to eavesdrop traffic on the network interface card (NIC) or switch level.

Audit:

Enter the following command:

```
lsqeth | grep isolation
```

Verify that the isolation parameter for all NICs is set to drop:

```
isolation : drop
```

Remediation:

Enter the following command to enable traffic isolation for all NICs:

```
echo drop >/sys/devices/qeth/<devno>/isolation
```

Default Value:

none

3.2 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.

3.2.1 Ensure source routed packets are not accepted (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 commands and verify output matches:

```
# sysctl net.ipv4.conf.all.accept_source_route

net.ipv4.conf.all.accept_source_route = 0
# sysctl net.ipv4.conf.default.accept_source_route

net.ipv4.conf.default.accept_source_route = 0
# grep "net\.ipv4\.conf\.all\.accept_source_route" /etc/sysctl.conf
/etc/sysctl.d/*

net.ipv4.conf.all.accept_source_route= 0
# grep "net\.ipv4\.conf\.default\.accept_source_route" /etc/sysctl.conf
/etc/sysctl.d/*

net.ipv4.conf.default.accept_source_route= 0
# sysctl net.ipv6.conf.all.accept_source_route

net.ipv6.conf.all.accept_source_route = 0
# sysctl net.ipv6.conf.default.accept_source_route

net.ipv6.conf.default.accept_source_route = 0
# grep "net\.ipv6\.conf\.all\.accept_source_route" /etc/sysctl.conf
/etc/sysctl.d/*

net.ipv4.conf.all.accept_source_route= 0
# grep "net\.ipv6\.conf\.default\.accept_source_route" /etc/sysctl.conf
/etc/sysctl.d/*

net.ipv6.conf.default.accept_source_route= 0
```

Remediation:

Set the following parameters in `/etc/sysctl.conf` or a `/etc/sysctl.d/*` file:

```
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
```

Run the following commands 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.ipv6.conf.all.accept_source_route=0
# sysctl -w net.ipv6.conf.default.accept_source_route=0
# sysctl -w net.ipv4.route.flush=1
# sysctl -w net.ipv6.route.flush=1
```

3.2.2 Ensure ICMP redirects are not accepted (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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. By setting `net.ipv4.conf.all.accept_redirects` and `net.ipv6.conf.all.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.

Rationale:

Attackers could use bogus ICMP redirect messages to maliciously alter the system routing tables and get them to send packets to incorrect networks and allow your system packets to be captured.

Audit:

Run the following commands and verify output matches:

```
# sysctl net.ipv4.conf.all.accept_redirects

net.ipv4.conf.all.accept_redirects = 0
# sysctl net.ipv4.conf.default.accept_redirects

net.ipv4.conf.default.accept_redirects = 0
# grep "net\.ipv4\.conf\.all\.accept_redirects" /etc/sysctl.conf
/etc/sysctl.d/*

net.ipv4.conf.all.accept_redirects= 0
# grep "net\.ipv4\.conf\.default\.accept_redirects" /etc/sysctl.conf
/etc/sysctl.d/*

net.ipv4.conf.default.accept_redirects= 0
# sysctl net.ipv6.conf.all.accept_redirects

net.ipv6.conf.all.accept_redirects = 0
# sysctl net.ipv6.conf.default.accept_redirects

net.ipv6.conf.default.accept_redirects = 0
# grep "net\.ipv6\.conf\.all\.accept_redirects" /etc/sysctl.conf
/etc/sysctl.d/*

net.ipv6.conf.all.accept_redirects= 0
# grep "net\.ipv6\.conf\.default\.accept_redirects" /etc/sysctl.conf
/etc/sysctl.d/*

net.ipv6.conf.default.accept_redirects= 0
```

Remediation:

Set the following parameters in `/etc/sysctl.conf` or a `/etc/sysctl.d/*` file:

```
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
```

Run the following commands 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.ipv6.conf.all.accept_redirects=0
# sysctl -w net.ipv6.conf.default.accept_redirects=0
# sysctl -w net.ipv4.route.flush=1
# sysctl -w net.ipv6.route.flush=1
```

3.2.3 Ensure secure ICMP redirects are not accepted (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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` to 0 protects the system from routing table updates by possibly compromised known gateways.

Audit:

Run the following commands and verify output matches:

```
# sysctl net.ipv4.conf.all.secure_redirects

net.ipv4.conf.all.secure_redirects = 0
# sysctl net.ipv4.conf.default.secure_redirects

net.ipv4.conf.default.secure_redirects = 0
# grep "net\.ipv4\.conf\.all\.secure_redirects" /etc/sysctl.conf
/etc/sysctl.d/*

net.ipv4.conf.all.secure_redirects= 0
# grep "net\.ipv4\.conf\.default\.secure_redirects" /etc/sysctl.conf
/etc/sysctl.d/*

net.ipv4.conf.default.secure_redirects= 0
```

Remediation:

Set the following parameters in `/etc/sysctl.conf` or a `/etc/sysctl.d/*` file:

```
net.ipv4.conf.all.secure_redirects = 0
net.ipv4.conf.default.secure_redirects = 0
```

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
```

3.2.4 Ensure suspicious packets are logged (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

When enabled, this feature logs packets with un-routable source addresses to the kernel log.

Rationale:

Enabling this feature and logging these packets allows an administrator to investigate the possibility that an attacker is sending spoofed packets to their system.

Audit:

Run the following commands and verify output matches:

```
# sysctl net.ipv4.conf.all.log_martians

net.ipv4.conf.all.log_martians = 1
# sysctl net.ipv4.conf.default.log_martians

net.ipv4.conf.default.log_martians = 1
# grep "net\.ipv4\.conf\.all\.log_martians" /etc/sysctl.conf /etc/sysctl.d/*

net.ipv4.conf.all.log_martians = 1
# grep "net\.ipv4\.conf\.default\.log_martians" /etc/sysctl.conf
/etc/sysctl.d/*

net.ipv4.conf.default.log_martians = 1
```

Remediation:

Set the following parameters in `/etc/sysctl.conf` or a `/etc/sysctl.d/*` file:

```
net.ipv4.conf.all.log_martians = 1
net.ipv4.conf.default.log_martians = 1
```

Run the following commands 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
```

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.			

3.2.5 Ensure broadcast ICMP requests are ignored (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 commands and verify output matches:

```
# sysctl net.ipv4.icmp_echo_ignore_broadcasts

net.ipv4.icmp_echo_ignore_broadcasts = 1
# grep -E -s "^s*net\.ipv4\.icmp_echo_ignore_broadcasts\s*=\s*0"
/etc/sysctl.conf /etc/sysctl.d/*.conf /usr/lib/sysctl.d/*.conf
/run/sysctl.d/*.conf

Nothing should be returned
```

Remediation:

Run the following command to restore the default parameters and set the active kernel parameters:

```
# grep -Els "^s*net\.ipv4\.icmp_echo_ignore_broadcasts\s*=\s*0"
/etc/sysctl.conf /etc/sysctl.d/*.conf /usr/lib/sysctl.d/*.conf
/run/sysctl.d/*.conf | while read filename; do sed -ri
"s/^s*(net\.ipv4\.icmp_echo_ignore_broadcasts\s*) (=) (\s*\S+\b).*$/#
*REMOVED* \1/" $filename; done; sysctl -w net.icmp_echo_ignore_broadcasts=1;
sysctl -w net.ipv4.route.flush=1
```

3.2.6 Ensure bogus ICMP responses are ignored (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Setting `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 commands and verify output matches:

```
# sysctl net.ipv4.icmp_ignore_bogus_error_responses

net.ipv4.icmp_ignore_bogus_error_responses = 1
# grep -E -s "^s*net\.ipv4\.icmp_ignore_bogus_error_responses\s*=\s*0"
/etc/sysctl.conf /etc/sysctl.d/*.conf /usr/lib/sysctl.d/*.conf
/run/sysctl.d/*.conf

Nothing should be returned
```

Remediation:

Run the following commands to restore the default parameters and set the active kernel parameters:

```
# grep -Els "^s*net\.ipv4\.icmp_ignore_bogus_error_responses\s*=\s*0"
/etc/sysctl.conf /etc/sysctl.d/*.conf /usr/lib/sysctl.d/*.conf
/run/sysctl.d/*.conf | while read filename; do sed -ri
"s/^\s*(net\.ipv4\.icmp_ignore_bogus_error_responses\s*) (=) (\s*\S+\b) .*$/#
*REMOVED* \1/" $filename; done; sysctl -w
net.ipv4.icmp_ignore_bogus_error_responses=1; sysctl -w
net.ipv4.route.flush=1
```

Default Value:

`net.ipv4.icmp_ignore_bogus_error_responses = 1`

3.2.7 Ensure reverse path filtering is enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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).

If asymmetrical routing may be in use, be sure that 'agreed to value' reflects this.

```
net.ipv4.conf.all.rp_filter = 1
net.ipv4.conf.default.rp_filter = 1
```

Rationale:

Setting these flags 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.

Audit:

Run the following commands and verify output matches:

```
# sysctl net.ipv4.conf.all.rp_filter

net.ipv4.conf.all.rp_filter = 1
# sysctl net.ipv4.conf.default.rp_filter

net.ipv4.conf.default.rp_filter = 1
grep "net.ipv4.conf.all.rp_filter"
    /etc/sysctl.conf /etc/sysctl.d/* | LC_ALL=C sort -V

/* must return nothing -or- must return: */
net.ipv4.conf.all.rp_filter = 1

/* Note: If the command returns multiple results,
the -last- entry must include the above line. */

grep "net.ipv4.conf.default.rp_filter"
    /etc/sysctl.conf /etc/sysctl.d/* | LC_ALL=C sort -V

/* must return nothing -or- must return: */
net.ipv4.conf.default.rp_filter = 1

/* Note: If the command returns multiple results,
the -last- entry must include the above line. */
```

Remediation:

Set the following parameters in /etc/sysctl.d/zzz-IBM.conf:

```
net.ipv4.conf.all.rp_filter = 1
net.ipv4.conf.default.rp_filter = 1
```

Revalidate that these are the last lines returned;

remove lines from other files if necessary.

Or simply eliminate those lines in -all-

/etc/sysctl.conf and /etc/sysctl.d/* files.

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
```

Default Value:

net.ipv4.conf.all.rp_filter = 1

net.ipv4.conf.default.rp_filter = 0

3.2.8 Ensure TCP SYN Cookies is enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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. SYN cookies allow the system to keep accepting valid connections, even if under a denial of service attack.

Audit:

Run the following commands and verify output matches:

```
# sysctl net.ipv4.tcp_syncookies

net.ipv4.tcp_syncookies = 1
# grep -E -r "^s*net\.ipv4\.tcp_syncookies\s*=\s*[02]*" /etc/sysctl.conf
/etc/sysctl.d/*.conf /usr/lib/sysctl.d/*.conf /run/sysctl.d/*.conf
```

Nothing should be returned

Remediation:

Run the following command to restore the default parameter and set the active kernel parameters:

```
# grep -Els "^s*net\.ipv4\.tcp_syncookies\s*=\s*[02]*" /etc/sysctl.conf
/etc/sysctl.d/*.conf /usr/lib/sysctl.d/*.conf /run/sysctl.d/*.conf | while
read filename; do sed -ri
"s/^s*(net\.ipv4\.tcp_syncookies\s*) (=) (\s*\S+\b).*$/# *REMOVED* \1/"
$filename; done; sysctl -w net.ipv4.tcp_syncookies=1; sysctl -w
net.ipv4.route.flush=1
```

3.2.9 Ensure IPv6 router advertisements are not accepted (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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.

Audit:

Run the following commands and verify output matches:

```
# sysctl net.ipv6.conf.all.accept_ra

net.ipv6.conf.all.accept_ra = 0
# sysctl net.ipv6.conf.default.accept_ra

net.ipv6.conf.default.accept_ra = 0
# grep "net\.ipv6\.conf\.all\.accept_ra" /etc/sysctl.conf /etc/sysctl.d/*

net.ipv6.conf.all.accept_ra = 0
# grep "net\.ipv6\.conf\.default\.accept_ra" /etc/sysctl.conf /etc/sysctl.d/*

net.ipv6.conf.default.accept_ra = 0
```

Remediation:

Set the following parameters in `/etc/sysctl.conf` or a `/etc/sysctl.d/*` file:

```
net.ipv6.conf.all.accept_ra = 0
net.ipv6.conf.default.accept_ra = 0
```

Run the following commands 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
```

3.3 Uncommon Network Protocols

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.

Note: This should not be considered a comprehensive list of uncommon network protocols, you may wish to consider additions to those listed here for your environment.

3.3.1 Ensure DCCP is disabled (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

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 commands and verify the output is as indicated:

```
# modprobe -n -v dccp
install /bin/true

# lsmod | grep dccp
<No output>
```

Remediation:

Edit or create a file in the `/etc/modprobe.d/` directory ending in `.conf`

Example: `vim /etc/modprobe.d/dccp.conf`

and add the following line:

```
install dccp /bin/true
```

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.			

3.3.2 Disable forwarding of packets (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Depending on its purpose, some system should never serve as a router and therefore, the forwarding of packets must be disabled. 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.

Setting these flags to 0 ensures that a system is not allowed to forward packets.

You can specify this TCP/IP stack setting either in the file `/etc/sysctl.conf`, or in some other control file in directory `/etc/sysctl.d`. Or you can just use the default (`ip_forward = 0`) on systems that are not routers, to not allow forwarding of packets.

Rationale:

Setting the flag to 0 ensures that a system with multiple interfaces, for example, a hard proxy, is unable to forward packets, and therefore, can never serve as a router.

Note: If a customer needs Linux systems to work as routers, indicate as an agreed value how to document which systems will be enabled to serve as a router.

Audit:

Run the listed commands and verify that the output is as required:

- Enter the following command:

```
sysctl net.ipv4.ip_forward
```

The output should return the following information:

```
net.ipv4.ip_forward = 0
```

- Enter the following command:

```
grep "net.ipv4.ip_forward" /etc/sysctl.conf /etc/sysctl.d/* | LC_ALL=C sort -V
```

The output should return the following information:

```
net.ipv4.ip_forward = 0
```

Note: If the resulting output returns multiple lines, this must be the **last** returned line.

Remediation:

Set the following parameter in file `/etc/sysctl.d/zzz-IBM.conf`:

```
net.ipv4.ip_forward = 0
```

Then rerun the commands in the **Audit Procedure** to verify the results. If required, remove the respective lines from the matched files.

Run the following commands to set the active kernel parameters:

```
# sysctl -w net.ipv4.ip_forward=0
# sysctl -w net.ipv4.route.flush=1
```

Default Value:

N/A

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.			

3.3.3 Ensure SCTP is disabled (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

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 commands and verify the output is as indicated:

```
# modprobe -n -v sctp
install /bin/true

# lsmod | grep sctp
<No output>
```

Remediation:

Edit or create a file in the `/etc/modprobe.d/` directory ending in `.conf`
Example: `vim /etc/modprobe.d/sctp.conf`
and add the following line:

```
install sctp /bin/true
```


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.			

3.3.4 Ensure RDS is disabled (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

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 commands and verify the output is as indicated:

```
# modprobe -n -v rds
install /bin/true

# lsmod | grep rds
<No output>
```

Remediation:

Edit or create a file in the `/etc/modprobe.d/` directory ending in `.conf`

Example: `vim /etc/modprobe.d/rds.conf`

and add the following line:

```
install rds /bin/true
```

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.			

3.3.5 Ensure TIPC is disabled (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

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 commands and verify the output is as indicated:

```
# modprobe -n -v tipc
install /bin/true

# lsmod | grep tipc
<No output>
```

Remediation:

Edit or create a file in the `/etc/modprobe.d/` directory ending in `.conf`

Example: `vim /etc/modprobe.d/tipc.conf`
and add the following line:

```
install tipc /bin/true
```

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.			

3.4 Firewall Configuration

A firewall Provides defense against external and internal threats by refusing unauthorized connections, to stop intrusion and provide a strong method of access control policy.

this section is intended only to ensure the resulting firewall rules are in place, not how they are configured.

3.4.1 Ensure Firewall software is installed

In order to configure Firewall protection for your system, a Firewall utility needs to be installed. Guidance has been included for the following firewall utilities.

- firewalld: Use the firewalld utility to configure less complex firewalls. The utility is easy to use and covers the typical use cases scenario.
- nftables: Use the nftables utility to set up complex firewalls, such as for a whole network.
- iptables: The iptables utility is deprecated in Fedora 28 derived Linux distributions. It is recommended that nftables be used instead. Guidance for iptables has been included for local sites that still require the use of iptables. This guidance may be removed in future releases of this Benchmark

Only one firewall utility should be installed and configured

3.4.1.1 Ensure a Firewall package is installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

A Firewall package should be selected. Most firewall configuration utilities operate as a front end to nftables or iptables.

Rationale:

A Firewall package is required for firewall management and configuration.

Audit:

Run **one** of the following commands to verify the Firewall package is installed:

For firewalld:

```
# rpm -q firewalld
```

For nftables:

```
# rpm -q nftables
```

For iptables:

```
# rpm -q iptables
```

Remediation:

Run **one** of the following commands to install a Firewall package.

For firewalld:

```
# dnf install firewalld
```

For nftables:

```
# dnf install nftables
```

For iptables:

```
# dnf install iptables
```

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 Configure firewalld

firewalld (Dynamic Firewall Manager) provides a dynamically managed firewall with support for network/firewall “zones” to assign a level of trust to a network and its associated connections, interfaces or sources. It has support for IPv4, IPv6, Ethernet bridges and also for IPSet firewall settings. There is a separation of the runtime and permanent configuration options. It also provides an interface for services or applications to add iptables, ip6tables and ebtables rules directly. This interface can also be used by advanced users.

Note: In the v0.6.0 release, firewalld gained support for using nftables as a firewall back-end. In CentOS Linux 8 firewalld utilizes nftables by default. Guidance within this section assumes that nftables are being used as the back-end of firewalld

3.4.2.1 Ensure firewalld service is enabled and running (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Ensure that the firewalld service is enabled to protect your system

Rationale:

firewalld (Dynamic Firewall Manager) tool provides a dynamically managed firewall. The tool enables network/firewall zones to define the trust level of network connections and/or interfaces. It has support both for IPv4 and IPv6 firewall settings. Also, it supports Ethernet bridges and allow you to separate between runtime and permanent configuration options. Finally, it supports an interface for services or applications to add firewall rules directly

Impact:

Changing firewall settings while connected over network can result in being locked out of the system.

Audit:

Run the following command to verify that firewalld is enabled:

```
# systemctl is-enabled firewalld  
enabled
```

Run the following command to verify that firewalld is running

```
# firewall-cmd --state  
running
```

Remediation:

Run the following command to enable and start firewalld

```
# systemctl --now enable firewalld
```

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.			

3.4.2.2 Ensure iptables is not enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

IPtables is an application that allows a system administrator to configure the IPv4 and IPv6 tables, chains and rules provided by the Linux kernel firewall.

IPtables is installed as a dependency with firewalld.

Rationale:

Running firewalld and IPtables concurrently may lead to conflict, therefore IPtables should be stopped and masked when using firewalld.

Audit:

Run the following command to verify that iptables is not running:

```
# systemctl status iptables
```

Output should include:

```
Loaded: disabled (/dev/null; bad)
Active: inactive (dead)
```

Run the following command to verify that iptables is not enabled:

```
# systemctl is-enabled iptables
```

Output should not read `enabled`

Remediation:

Run the following command to stop and mask iptables

```
systemctl --now mask iptables
```

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.			

3.4.2.3 Ensure nftables is not enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

nftables is a subsystem of the Linux kernel providing filtering and classification of network packets/datagrams/frames and is the successor to iptables.

nftables are installed as a dependency with firewalld.

Rationale:

Running firewalld and nftables concurrently may lead to conflict, therefore nftables should be stopped and masked when using firewalld.

Audit:

Run the following command to verify that nftables is not enabled:

```
# systemctl is-enabled nftables  
  
(disabled|masked)
```

Run the following command to verify that nftables is not running:

```
# systemctl status nftables
```

Output should include:

```
Loaded: masked (/dev/null; bad)  
Active: inactive (dead)
```

Remediation:

Run the following command to mask and stop nftables

```
systemctl --now mask nftables
```

Additional Information:

firewalld is dependent on nftables. nftables should be stopped and disabled.

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.			

3.4.2.4 Ensure default zone is set (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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.

That means that if there is 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.

Rationale:

Because the default zone is the zone that is used for everything that is not explicitly bound/assigned to another zone, it is important for the default zone to set

Audit:

Run the following command and verify that the default zone adheres to company policy:

```
# firewall-cmd --get-default-zone
```

Remediation:

Run the following command to set the default zone:

```
# firewall-cmd --set-default-zone=<NAME_OF_ZONE>
```

Example:

```
# firewall-cmd --set-default-zone=public
```

References:

1. <https://firewalld.org/documentation>
2. <https://firewalld.org/documentation/man-pages/firewalld.zone>

Additional Information:

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

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.			

3.4.2.5 Ensure network interfaces are assigned to appropriate zone (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

firewall zones define the trust level of network connections or interfaces.

Rationale:

A network interface not assigned to the appropriate zone can allow unexpected or undesired network traffic to be accepted on the interface

Impact:

Changing firewall settings while connected over network can result in being locked out of the system.

Audit:

Run the following command, and verify that the interface(s) follow site policy for zone assignment

```
# nmcli -t connection show | awk -F: '{if($4){print $4}}' | while read INT;
do firewall-cmd --get-active-zones | grep -B1 $INT; done
```

Remediation:

Run the following command to assign an interface to the appropriate zone.

```
# firewall-cmd --zone=<Zone NAME> --change-interface=<INTERFACE NAME>
```

Example:

```
# firewall-cmd --zone=customzone --change-interface=eth0
```

Default Value:

If the zone is not set in the configuration file, the interfaces will be assigned to the default zone defined in the firewalld configuration

References:

1. <https://firewalld.org/documentation/zone/connections-interfaces-and-sources.html>

Additional Information:

The firewall in the Linux kernel is not able to handle network connections with the name shown by NetworkManager, it can only handle the network interfaces used by the connection. Because of this NetworkManager tells firewalld to assign the network interface that is used for this connection to the zone defined in the configuration of that connection. This assignment happens before the interface is used. The configuration of the connection can either be the NetworkManager configuration or also an ifcfg for example. If the zone is not set in the configuration file, the interfaces will be assigned to the default zone defined in the firewalld configuration. If a connection has more than one interface, all of them will be supplied to firewalld. Also changes in the names of interfaces will be handled by NetworkManager and supplied to firewalld.

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.			

3.4.2.6 *Ensure unnecessary services and ports are not accepted (Manual)*

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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.

By setting the target to ACCEPT, you accept all incoming packets except those disabled by a specific rule.

If you set the target to REJECT or DROP, you disable all incoming packets except those that you have allowed in specific rules. When packets are rejected, the source machine is informed about the rejection, while there is no information sent when the packets are dropped.

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.

```
# firewall-cmd --get-active-zones | awk '!/:/ {print $1}' | while read ZN; do  
firewall-cmd --list-all --zone=$ZN; done
```

Remediation:

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. firewall.service(5)
2. https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html/securing_networks/using-and-configuring-firewalls_securing-networks

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.			

3.4.3 Configure nftables

nftables is a subsystem of the Linux kernel providing filtering and classification of network packets/datagrams/frames and is the successor to iptables. The biggest change with the successor nftables is its simplicity. With iptables, we have to configure every single rule and use the syntax which can be compared with normal commands. With nftables, the simpler syntax, much like BPF (Berkely Packet Filter) means shorter lines and less repetition. Support for nftables should also be compiled into the kernel, together with the related nftables modules. Please ensure that your kernel supports `nf_tables` before choosing this option.

Note: This section broadly assumes starting with an empty nftables firewall ruleset (established by flushing the rules with `nft flush ruleset`). Remediation steps included only affect the live system, you will also need to configure your default firewall configuration to apply on boot. 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.

The following will implement the firewall rules of this section and 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) needs to be updated to only allow systems requiring ssh connectivity to connect, as per site policy.

Save the script bellow as `/etc/nftables/nftables.rules`

```
#!/sbin/nft -f

# 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 {
    # Base chain for input hook named input (Filters inbound network packets)
    chain input {
        type filter hook input priority 0; policy drop;

        # Ensure loopback traffic is configured
        iif "lo" accept
        ip saddr 127.0.0.0/8 counter packets 0 bytes 0 drop
        ip6 saddr ::1 counter packets 0 bytes 0 drop

        # Ensure established connections are configured
        ip protocol tcp ct state established accept
        ip protocol udp ct state established accept
        ip protocol icmp ct state established accept

        # Accept port 22(SSH) traffic from anywhere
        tcp dport ssh accept

        # Accept ICMP and IGMP from anywhere
        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
        ip protocol igmp accept
    }

    # Base chain for hook forward named forward (Filters forwarded network
packets)
    chain forward {
        type filter hook forward priority 0; policy drop;
    }

    # Base chain for hook output named output (Filters outbound network packets)
    chain output {
        type filter hook output priority 0; policy drop;
        # Ensure outbound and established connections are configured
        ip protocol tcp ct state established,related,new accept
        ip protocol udp ct state established,related,new accept
        ip protocol icmp ct state established,related,new accept
    }
}
}
```

Run the following command to load the file into nftables

```
# nft -f /etc/nftables/nftables.rules
```

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"
```


3.4.3.1 Ensure iptables are flushed (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

nftables is a replacement for iptables, ip6tables, ebtables and arptables

Rationale:

It is possible to mix iptables and nftables. However, this increases complexity and also the chance to introduce errors. For simplicity flush out all iptables rules, and ensure it is not loaded

Audit:

Run the following commands to ensure no iptables rules exist

For iptables:

```
# iptables -L
```

No rules should be returned

For ip6tables:

```
# ip6tables -L
```

No rules should be returned

Remediation:

Run the following commands to flush iptables:

For iptables:

```
# iptables -F
```

For ip6tables

```
# ip6tables -F
```

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.			

3.4.3.2 Ensure a table exists (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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:

nftables doesn't have any default tables. Without a table being build, nftables will not filter network traffic.

Impact:

Adding rules to a running nftables 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:
example:

```
table inet filter
```

Remediation:

Run the following command to create a table in nftables

```
# nft create table inet <table name>
```

Example:

```
# nft create table inet filter
```

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.			

3.4.3.3 Ensure base chains exist (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 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:

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;
```

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 \; }
```

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.			

3.4.3.4 Ensure loopback traffic is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 a 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 commands to verify that the loopback interface is configured:

```
# nft list ruleset | awk '/hook input/,/}/' | grep 'iif "lo" accept'

iif "lo" accept
# nft list ruleset | awk '/hook input/,/}/' | grep 'ip saddr'

ip saddr 127.0.0.0/8 counter packets 0 bytes 0 drop
# nft list ruleset | awk '/hook input/,/}/' | grep 'ip6 saddr'

ip6 saddr ::1 counter packets 0 bytes 0 drop
```

Remediation:

Run the following commands to implement the loopback rules:

```
# nft add rule inet filter input iif lo accept
# nft create rule inet filter input ip saddr 127.0.0.0/8 counter drop
# nft add rule inet filter input ip6 saddr ::1 counter drop
```

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.			

3.4.3.5 Ensure outbound and established connections are configured (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Configure the firewall rules for new outbound, and established connections

Rationale:

If rules are not in place for new outbound, and established connections all packets will be dropped by the default policy preventing network usage.

Audit:

Run the following commands and verify all rules for established incoming connections match site policy: site policy:

```
# nft list ruleset | awk '/hook input/,/}/' | grep -E 'ip protocol  
(tcp|udp|icmp) 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
```

Run the following command and verify all rules for new and established outbound connections match site policy

```
# nft list ruleset | awk '/hook output/,/}/' | grep -E 'ip protocol  
(tcp|udp|icmp) ct state'
```

Output should be similar to:

```
ip protocol tcp ct state established,related,new accept  
ip protocol udp ct state established,related,new accept  
ip protocol icmp ct state established,related,new accept
```

Remediation:

Configure nftables in accordance with site policy. The following commands will implement a policy to allow all outbound connections and all established connections:

```
# nft add rule inet filter input ip protocol tcp ct state established accept
# nft add rule inet filter input ip protocol udp ct state established accept
# nft add rule inet filter input ip protocol icmp ct state established accept

# nft add rule inet filter output ip protocol tcp ct state
new,related,established accept

# nft add rule inet filter output ip protocol udp ct state
new,related,established accept

# nft add rule inet filter output ip protocol icmp ct state
new,related,established accept
```

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.			

3.4.3.6 Ensure default deny firewall policy (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 transversing the network stack.

It is easier to white list acceptable usage than to black list unacceptable usage.

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:

Run the following commands and verify that base chains contain a policy of `DROP`.

```
# nft list ruleset | grep 'hook input'

type filter hook input priority 0; policy drop;
# nft list ruleset | grep 'hook forward'

type filter hook forward priority 0; policy drop;
# nft list ruleset | grep 'hook output'

type filter hook output priority 0; policy drop;
```

Remediation:

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 \; }
# nft chain inet filter output { policy drop \; }
```

Default Value:

accept

References:

1. Manual Page nft

Additional Information:

Changing firewall settings while connected over network can result in being locked out of the system.

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.			

3.4.3.7 Ensure nftables service is enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The nftables service allows for the loading of nftables rulesets during boot, or starting of the nftables service

Rationale:

The nftables service restores the nftables rules from the rules files referenced in the `/etc/sysconfig/nftables.conf` file during boot or the starting of the nftables service

Audit:

Run the following command and verify that the nftables service is enabled:

```
# systemctl is-enabled nftables  
  
enabled
```

Remediation:

Run the following command to enable the nftables service:

```
# systemctl --now enable nftables
```

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.			

3.4.3.8 Ensure nftables rules are permanent (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

nftables is a subsystem of the Linux kernel providing filtering and classification of network packets/datagrams/frames.

The nftables service reads the `/etc/sysconfig/nftables.conf` file for a nftables file or files to include in the nftables ruleset.

A nftables ruleset containing the input, forward, and output base chains allow network traffic to be filtered.

Rationale:

Changes made to nftables ruleset only affect the live system, you will also need to configure the nftables ruleset to apply on boot.

Audit:

Run the following commands to verify that input, forward, and output base chains are configured to be applied to a nftables ruleset on boot:

Run the following command to verify the input base chain:

```
# [[ -n $(grep -E "^s*include" /etc/sysconfig/nftables.conf) ]] && awk
'/hook input/,/}/' $(awk '$1 ~ /^s*include/ { gsub("\"", "", $2); print $2 }'
/etc/sysconfig/nftables.conf)
```

Output should be similar to:

```
type filter hook input priority 0; policy drop;

# Ensure loopback traffic is configured
iif "lo" accept
ip saddr 127.0.0.0/8 counter packets 0 bytes 0 drop
ip6 saddr ::1 counter packets 0 bytes 0 drop

# Ensure established connections are configured
ip protocol tcp ct state established accept
ip protocol udp ct state established accept
ip protocol icmp ct state established accept

# Accept port 22 (SSH) traffic from anywhere
tcp dport ssh accept

# Accept ICMP and IGMP from anywhere
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
```

Note: Review the input base chain to ensure that it follows local site policy

Run the following command to verify the forward base chain:

```
# [[ -n $(grep -E "^s*include" /etc/sysconfig/nftables.conf) ]] && awk
'/hook forward/,/}/' $(awk '$1 ~ /^s*include/ { gsub("\"", "", $2); print $2 }'
/etc/sysconfig/nftables.conf)
```

Output should be similar to:

```
# Base chain for hook forward named forward (Filters forwarded
network packets)
chain forward {
    type filter hook forward priority 0; policy drop;
}
```

Note: Review the forward base chain to ensure that it follows local site policy.

Run the following command to verify the output base chain:

```
# [[ -n $(grep -E "^s*include" /etc/sysconfig/nftables.conf) ]] && awk
'/hook output/,/}/' $(awk '$1 ~ /^s*include/ { gsub("\"", "", $2); print $2 }'
/etc/sysconfig/nftables.conf)
```

Output should be similar to:

```
# Base chain for hook output named output (Filters outbound network
packets)
chain output {
    type filter hook output priority 0; policy drop;
    # Ensure outbound and established connections are configured
    ip protocol tcp ct state established,related,new accept
    ip protocol tcp ct state established,related,new accept
    ip protocol udp ct state established,related,new accept
    ip protocol icmp ct state established,related,new accept
}
```

Note: Review the output base chain to ensure that it follows local site policy.

Remediation:

Edit the `/etc/sysconfig/nftables.conf` file and un-comment or add a line with `include` `<Absolute path to nftables rules file>` for each nftables file you want included in the nftables ruleset on boot
example:

```
# vi /etc/sysconfig/nftables.conf
```

Add the line:

```
include "/etc/nftables/nftables.rules"
```

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.			

3.4.4 Configure iptables

IPtables is an application that allows a system administrator to configure the IPv4 and IPv6 tables, chains and rules provided by the Linux kernel firewall. While several methods of configuration exist this section is intended only to ensure the resulting IPtables rules are in place, not how they are configured. If IPv6 is in use in your environment, similar settings should be applied to the IP6tables as well.

3.4.4.1 Configure software

This section provides guidance for installing, enabling, removing, and disabling software packages necessary for using IPTables as the method for configuring and maintaining a Host Based Firewall on the system.

Note: Using more than one method to configure and maintain a Host Based Firewall can cause unexpected results. If FirewallD or nftables are being used for configuration and maintenance, this guidance should be based over and the guidance in their respective section followed.

3.4.4.1.1 Ensure iptables packages are installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

iptables is a utility program that allows a system administrator to configure the tables provided by the Linux kernel firewall, implemented as different Netfilter modules, and the chains and rules it stores. Different kernel modules and programs are used for different protocols; iptables applies to IPv4, ip6tables to IPv6, arptables to ARP, and ebtables to Ethernet frames.

Rationale:

A method of configuring and maintaining firewall rules is necessary to configure a Host Based Firewall.

Audit:

Run the following command to verify that `iptables` and `iptables-services` are installed:

```
rpm -q iptables iptables-services  
  
iptables-<version>  
iptables-services-<version>
```

Remediation:

Run the following command to install `iptables` and `iptables-services`

```
# yum install iptables iptables-services
```

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.			

3.4.4.1.2 Ensure nftables is not installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

nftables is a subsystem of the Linux kernel providing filtering and classification of network packets/datagrams/frames and is the successor to iptables.

Rationale:

Running both iptables and nftables may lead to conflict.

Audit:

Run the following command to verify that nftables is not installed:

```
# rpm -q nftables  
package nftables is not installed
```

Remediation:

Run the following command to remove nftables:

```
# yum remove nftables
```

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.			

3.4.4.1.3 Ensure firewalld is not installed or stopped and masked (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

firewalld (Dynamic Firewall Manager) provides a dynamically managed firewall with support for network/firewall “zones” to assign a level of trust to a network and its associated connections, interfaces or sources. It has support for IPv4, IPv6, Ethernet bridges and also for IPSet firewall settings. There is a separation of the runtime and permanent configuration options.

Rationale:

Running `iptables.service` and/or `ip6tables.service` with `firewalld.service` may lead to conflict and unexpected results.

Audit:

Run the following command to verify that `firewalld` is not installed:

```
# rpm -q firewalld  
  
package firewalld is not installed
```

OR

Run the following commands to verify that `firewalld` is stopped and masked

```
# systemctl status firewalld | grep "Active: " | grep -v "active (running) "  
  
No output should be returned  
# systemctl is-enabled firewalld  
  
masked
```

Remediation:

Run the following command to remove `firewalld`

```
# yum remove firewalld
```

OR

Run the following command to stop and mask `firewalld`

```
# systemctl --now mask firewalld
```

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.			

3.4.4.2 Configure IPv4 iptables

Iptables is used to set up, maintain, and inspect the tables of IP packet filter rules in the Linux kernel. Several different tables may be defined. Each table contains a number of built-in chains and may also contain user-defined chains.

Each chain is a list of rules which can match a set of packets. Each rule specifies what to do with a packet that matches. This is called a 'target', which may be a jump to a user-defined chain in the same table.

Note: This section broadly assumes starting with an empty IPtables firewall ruleset (established by flushing the rules with `iptables -F`). Remediation steps included only affect the live system, you will also need to configure your default firewall configuration to apply on boot. 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. The following script will implement the firewall rules of this section and open port 22(ssh) from anywhere. This needs to be updated to only allow systems requiring ssh connectivity to connect as per site policy.

```
#!/bin/bash

# Flush IPtables rules
iptables -F

# Ensure default deny firewall policy
iptables -P INPUT DROP
iptables -P OUTPUT DROP
iptables -P FORWARD DROP

# Ensure loopback traffic is configured
iptables -A INPUT -i lo -j ACCEPT
iptables -A OUTPUT -o lo -j ACCEPT
iptables -A INPUT -s 127.0.0.0/8 -j DROP

# Ensure outbound and established connections are configured
iptables -A OUTPUT -p tcp -m state --state NEW,ESTABLISHED -j ACCEPT
iptables -A OUTPUT -p udp -m state --state NEW,ESTABLISHED -j ACCEPT
iptables -A OUTPUT -p icmp -m state --state NEW,ESTABLISHED -j ACCEPT
iptables -A INPUT -p tcp -m state --state ESTABLISHED -j ACCEPT
iptables -A INPUT -p udp -m state --state ESTABLISHED -j ACCEPT
iptables -A INPUT -p icmp -m state --state ESTABLISHED -j ACCEPT

# Open inbound ssh(tcp port 22) connections
iptables -A INPUT -p tcp --dport 22 -m state --state NEW -j ACCEPT
```

3.4.4.2.1 Ensure default deny firewall policy (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

A default deny all policy on connections ensures that any unconfigured network usage will be rejected.

Rationale:

With a default accept policy the firewall will accept any packet that is not configured to be denied. It is easier to white list acceptable usage than to black list unacceptable usage.

Audit:

Run the following command and verify that the policy for the `INPUT` , `OUTPUT` , and `FORWARD` chains is `DROP` or `REJECT` :

```
# iptables -L
Chain INPUT (policy DROP)
Chain FORWARD (policy DROP)
Chain OUTPUT (policy DROP)
```

Remediation:

Run the following commands to implement a default DROP policy:

```
# iptables -P INPUT DROP
# iptables -P OUTPUT DROP
# iptables -P FORWARD DROP
```

Additional Information:

Changing firewall settings while connected over network can result in being locked out of the system.

Remediation will only affect the active system firewall, be sure to configure the default policy in your firewall management to apply on boot as well.

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.			

3.4.4.2.2 Ensure loopback traffic is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Configure the loopback interface to accept traffic. Configure all other interfaces to deny traffic to the loopback network (127.0.0.0/8).

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 (127.0.0.0/8) traffic should be seen, all other interfaces should ignore traffic on this network as an anti-spoofing measure.

Audit:

Run the following commands and verify output includes the listed rules in order (packet and byte counts may differ):

```
# iptables -L INPUT -v -n
Chain INPUT (policy DROP 0 packets, 0 bytes)
  pkts bytes target     prot opt in     out     source
destination
    0    0 ACCEPT     all  --  lo      *       0.0.0.0/0           0.0.0.0/0
    0    0 DROP       all  --  *       *       127.0.0.0/8         0.0.0.0/0

# iptables -L OUTPUT -v -n
Chain OUTPUT (policy DROP 0 packets, 0 bytes)
  pkts bytes target     prot opt in     out     source
destination
    0    0 ACCEPT     all  --  *       lo      0.0.0.0/0           0.0.0.0/0
```

Remediation:

Run the following commands to implement the loopback rules:

```
# iptables -A INPUT -i lo -j ACCEPT
# iptables -A OUTPUT -o lo -j ACCEPT
# iptables -A INPUT -s 127.0.0.0/8 -j DROP
```

Additional Information:

Changing firewall settings while connected over network can result in being locked out of the system.

Remediation will only affect the active system firewall, be sure to configure the default policy in your firewall management to apply on boot as well.

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.			

3.4.4.2.3 Ensure outbound and established connections are configured (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Configure the firewall rules for new outbound, and established connections.

Rationale:

If rules are not in place for new outbound, and established connections all packets will be dropped by the default policy preventing network usage.

Audit:

Run the following command and verify all rules for new outbound, and established connections match site policy:

```
# iptables -L -v -n
```

Remediation:

Configure iptables in accordance with site policy. The following commands will implement a policy to allow all outbound connections and all established connections:

```
# iptables -A OUTPUT -p tcp -m state --state NEW,ESTABLISHED -j ACCEPT
# iptables -A OUTPUT -p udp -m state --state NEW,ESTABLISHED -j ACCEPT
# iptables -A OUTPUT -p icmp -m state --state NEW,ESTABLISHED -j ACCEPT
# iptables -A INPUT -p tcp -m state --state ESTABLISHED -j ACCEPT
# iptables -A INPUT -p udp -m state --state ESTABLISHED -j ACCEPT
# iptables -A INPUT -p icmp -m state --state ESTABLISHED -j ACCEPT
```

Additional Information:

Changing firewall settings while connected over network can result in being locked out of the system.

Remediation will only affect the active system firewall, be sure to configure the default policy in your firewall management to apply on boot as well.

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.			

3.4.4.2.4 Ensure firewall rules exist for all open ports (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Any ports that have been opened on non-loopback addresses need firewall rules to govern traffic.

Rationale:

Without a firewall rule configured for open ports default firewall policy will drop all packets to these ports.

Audit:

Run the following command to determine open ports:

```
# ss -4tuln
```

Netid	State	Recv-Q	Send-Q	Local Address:Port	Peer
Address:Port					
udp	UNCONN	0	0	*:68	
:					
udp	UNCONN	0	0	*:123	
:					
tcp	LISTEN	0	128	*:22	
:					

Run the following command to determine firewall rules:

```
# iptables -L INPUT -v -n
```

Chain INPUT (policy DROP 0 packets, 0 bytes)									
	pkts	bytes	target	prot	opt	in	out	source	
destination									
	0	0	ACCEPT	all	--	lo	*	0.0.0.0/0	0.0.0.0/0
	0	0	DROP	all	--	*	*	127.0.0.0/8	0.0.0.0/0
	0	0	ACCEPT	tcp	--	*	*	0.0.0.0/0	0.0.0.0/0
tcp dpt:22 state NEW									

Verify all open ports listening on non-localhost addresses have at least one firewall rule. The last line identified by the "tcp dpt:22 state NEW" identifies it as a firewall rule for new connections on tcp port 22.

Remediation:

For each port identified in the audit which does not have a firewall rule establish a proper rule for accepting inbound connections:

```
# iptables -A INPUT -p <protocol> --dport <port> -m state --state NEW -j ACCEPT
```

Additional Information:

Changing firewall settings while connected over network can result in being locked out of the system.

Remediation will only affect the active system firewall, be sure to configure the default policy in your firewall management to apply on boot as well.

The remediation command opens up the port to traffic from all sources. Consult iptables documentation and set any restrictions in compliance with site policy.

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.			

3.4.4.3 Configure IPv6 ip6tables

Ip6tables is used to set up, maintain, and inspect the tables of IPv6 packet filter rules in the Linux kernel. Several different tables may be defined. Each table contains a number of built-in chains and may also contain user-defined chains. Each chain is a list of rules which can match a set of packets. Each rule specifies what to do with a packet that matches. This is called a 'target', which may be a jump to a user-defined chain in the same table.

If IPv6 is enabled on the system, the ip6tables should be configured.

Note: This section broadly assumes starting with an empty ip6tables firewall ruleset (established by flushing the rules with `ip6tables -F`). Remediation steps included only affect the live system, you will also need to configure your default firewall configuration to apply on boot. 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. The following script will implement the firewall rules of this section and open port 22(ssh) from anywhere. This needs to be updated to only allow systems requiring ssh connectivity to connect as per site policy.

```
#!/bin/bash

# Flush ip6tables rules
ip6tables -F

# Ensure default deny firewall policy
ip6tables -P INPUT DROP
ip6tables -P OUTPUT DROP
ip6tables -P FORWARD DROP

# Ensure loopback traffic is configured
ip6tables -A INPUT -i lo -j ACCEPT
ip6tables -A OUTPUT -o lo -j ACCEPT
ip6tables -A INPUT -s ::1 -j DROP

# Ensure outbound and established connections are configured
ip6tables -A OUTPUT -p tcp -m state --state NEW,ESTABLISHED -j ACCEPT
ip6tables -A OUTPUT -p udp -m state --state NEW,ESTABLISHED -j ACCEPT
ip6tables -A OUTPUT -p icmp -m state --state NEW,ESTABLISHED -j ACCEPT
ip6tables -A INPUT -p tcp -m state --state ESTABLISHED -j ACCEPT
ip6tables -A INPUT -p udp -m state --state ESTABLISHED -j ACCEPT
ip6tables -A INPUT -p icmp -m state --state ESTABLISHED -j ACCEPT

# Open inbound ssh(tcp port 22) connections
ip6tables -A INPUT -p tcp --dport 22 -m state --state NEW -j ACCEPT
```


3.4.4.3.1 Ensure IPv6 default deny firewall policy (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

A default deny all policy on connections ensures that any unconfigured network usage will be rejected.

Rationale:

With a default accept policy the firewall will accept any packet that is not configured to be denied. It is easier to white list acceptable usage than to black list unacceptable usage.

Audit:

If IPv6 is enabled on the system

Run the following command and verify that the policy for the INPUT, OUTPUT, and FORWARD chains is DROP or REJECT:

```
# ip6tables -L
Chain INPUT (policy DROP)
Chain FORWARD (policy DROP)
Chain OUTPUT (policy DROP)
```

Remediation:

Run the following commands to implement a default DROP policy:

```
# ip6tables -P INPUT DROP
# ip6tables -P OUTPUT DROP
# ip6tables -P FORWARD DROP
```

Additional Information:

Changing firewall settings while connected over network can result in being locked out of the system.

Remediation will only affect the active system firewall, be sure to configure the default policy in your firewall management to apply on boot as well.

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.			

3.4.4.3.2 Ensure IPv6 loopback traffic is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Configure the loopback interface to accept traffic. Configure all other interfaces to deny traffic to the loopback network (::1).

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 (::1) traffic should be seen, all other interfaces should ignore traffic on this network as an anti-spoofing measure.

Audit:

If IPv6 is enabled on the system

Run the following commands and verify output includes the listed rules in order (packet and byte counts may differ):

```
# ip6tables -L INPUT -v -n
Chain INPUT (policy DROP 0 packets, 0 bytes)
pkts bytes target      prot opt in      out     source
destination
    0     0 ACCEPT      all  *  lo      *      ::/0
    0     0 DROP        all  *  *       *      :::1

# ip6tables -L OUTPUT -v -n
Chain OUTPUT (policy DROP 0 packets, 0 bytes)
pkts bytes target      prot opt in      out     source
destination
    0     0 ACCEPT      all  *  *       lo     ::/0
```

Remediation:

Run the following commands to implement the loopback rules:

```
# ip6tables -A INPUT -i lo -j ACCEPT
# ip6tables -A OUTPUT -o lo -j ACCEPT
# ip6tables -A INPUT -s :::1 -j DROP
```

Additional Information:

Changing firewall settings while connected over network can result in being locked out of the system.

Remediation will only affect the active system firewall, be sure to configure the default policy in your firewall management to apply on boot as well.

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.			

3.4.4.3.3 Ensure IPv6 outbound and established connections are configured (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Configure the firewall rules for new outbound, and established IPv6 connections.

Rationale:

If rules are not in place for new outbound, and established connections all packets will be dropped by the default policy preventing network usage.

Audit:

If IPv6 is enabled on the system

Run the following command and verify all rules for new outbound, and established connections match site policy:

```
# ip6tables -L -v -n
```

Remediation:

Configure iptables in accordance with site policy. The following commands will implement a policy to allow all outbound connections and all established connections:

```
# ip6tables -A OUTPUT -p tcp -m state --state NEW,ESTABLISHED -j ACCEPT
# ip6tables -A OUTPUT -p udp -m state --state NEW,ESTABLISHED -j ACCEPT
# ip6tables -A OUTPUT -p icmp -m state --state NEW,ESTABLISHED -j ACCEPT
# ip6tables -A INPUT -p tcp -m state --state ESTABLISHED -j ACCEPT
# ip6tables -A INPUT -p udp -m state --state ESTABLISHED -j ACCEPT
# ip6tables -A INPUT -p icmp -m state --state ESTABLISHED -j ACCEPT
```

Additional Information:

Changing firewall settings while connected over network can result in being locked out of the system.

Remediation will only affect the active system firewall, be sure to configure the default policy in your firewall management to apply on boot as well.

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.			

3.4.4.3.4 Ensure IPv6 firewall rules exist for all open ports (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Any ports that have been opened on non-loopback addresses need firewall rules to govern traffic.

Rationale:

Without a firewall rule configured for open ports default firewall policy will drop all packets to these ports.

Audit:

If IPv6 is enabled on the system

Run the following command to determine open ports:

```
# ss -6tuln
```

Netid	State	Recv-Q	Send-Q	Local Address:Port	Peer
Address:Port					
udp	UNCONN	0	0	:::1:123	
:::*					
udp	UNCONN	0	0	:::123	
:::*					
tcp	LISTEN	0	128	:::22	
:::*					
tcp	LISTEN	0	20	:::1:25	
:::*					

Run the following command to determine firewall rules:

```
# ip6tables -L INPUT -v -n
```

Chain INPUT (policy DROP 0 packets, 0 bytes)							
	pkts	bytes	target	prot	opt	in	out
destination							source
	0	0	ACCEPT	all		lo	:::/0
	0	0	DROP	all		*	:::1
	0	0	ACCEPT	tcp		*	:::/0
tcp dpt:22 state NEW							

Verify all open ports listening on non-localhost addresses have at least one firewall rule. The last line identified by the "tcp dpt:22 state NEW" identifies it as a firewall rule for new connections on tcp port 22.

Remediation:

For each port identified in the audit which does not have a firewall rule establish a proper rule for accepting inbound connections:

```
# iptables -A INPUT -p <protocol> --dport <port> -m state --state NEW -j ACCEPT
```

Additional Information:

Changing firewall settings while connected over network can result in being locked out of the system.

Remediation will only affect the active system firewall, be sure to configure the default policy in your firewall management to apply on boot as well.

The remediation command opens up the port to traffic from all sources. Consult iptables documentation and set any restrictions in compliance with site policy.

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.			

3.5 Ensure wireless interfaces are disabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 2 - Workstation

Description:

Wireless networking is used when wired networks are unavailable. CentOS Linux contains a wireless tool kit to allow system administrators to configure and use wireless networks.

Rationale:

If wireless is not to be used, wireless devices can be disabled to reduce the potential attack surface.

Impact:

Many if not all laptop workstations and some desktop workstations will connect via wireless requiring these interfaces be enabled.

Audit:

Run the following command to verify no wireless interfaces are active on the system:

```
# nmcli radio all
```

Output should look like:

WIFI-HW	WIFI	WWAN-HW	WWAN
enabled	disabled	enabled	disabled

Remediation:

Run the following command to disable any wireless interfaces:

```
# nmcli radio all off
```

Disable any wireless interfaces in your network configuration.

References:

1. nmcli(1) - Linux man page

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>15.4 Disable Wireless Access on Devices if Not Required</u> Disable wireless access on devices that do not have a business purpose for wireless access.			●
v7	<u>15.5 Limit Wireless Access on Client Devices</u> Configure wireless access on client machines that do have an essential wireless business purpose, to allow access only to authorized wireless networks and to restrict access to other wireless networks.			●

3.6 Disable IPv6 (Manual)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Although IPv6 has many advantages over IPv4, not all organizations have IPv6 or dual stack configurations implemented.

Rationale:

If IPv6 or dual stack is not to be used, it is recommended that IPv6 be disabled to reduce the attack surface of the system.

Audit:

Run the following command to verify

```
# grep -E "^\s*kernelopts=(\S+\s+)*ipv6\.disable=1\b\s*(\S+\s+)*$" /boot/grub2/grubenv
```

Output should include

```
ipv6.disable=1
```

Remediation:

Edit `/etc/default/grub` and add `ipv6.disable=1` to the `GRUB_CMDLINE_LINUX` parameters:

```
GRUB_CMDLINE_LINUX="ipv6.disable=1"
```

Run the following command to update the `grub2` configuration:

```
# grub2-mkconfig -o /boot/grub2/grub.cfg
```

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.			

3.7 Choose hardware-accelerated cipher suites for applications using SSL/TLS (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

For every application that uses SSL/TLS ensure to use cipher suites only that are accelerated by cryptographic hardware.

Rationale:

Hardware accelerated cryptographic algorithms can exploit hardware-specific features for enhanced performance and security.

Audit:

Auditing of this recommendation is application-specific.

Note: Applications that obey global operating system settings can be configured in `/etc/crypto-policies/config`. See also 1.10/1.11.

Remediation:

Define adequate cipher suites in application specific configurations and specify a minimum level of safe cipher suites in `/etc/crypto-policies/config`

Default Value:

n/a

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 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`).

4.1 Configure System Accounting (auditd)

System auditing, through `auditd`, allows system administrators to monitor their systems such that they can detect unauthorized access or modification of data. By default, `auditd` audits SELinux AVC denials, system logins, account modifications, and authentication events. Events will be logged to `/var/log/audit/audit.log`. The recording of these events will use a modest amount of disk space on a system. If significantly more events are captured, additional on system or off system storage may need to be allocated.

The recommendations in this section implement an audit policy that produces large quantities of logged data. In some environments it can be challenging to store or process these logs and as such they are marked as Level 2 for both Servers and Workstations.

Note: 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.

Note: Systems may have been customized to change the default `UID_MIN`. To confirm the `UID_MIN` for your system, run the following command:

```
awk '/^\s*UID_MIN/{print $2}' /etc/login.defs
```

If your systems' `UID_MIN` is not 1000, replace `audit>=1000` with `audit>=<UID_MIN for your system>` in the Audit and Remediation procedures.

Note: Once all configuration changes have been made to `/etc/audit/rules.d/audit.rules`, the `auditd` configuration must be reloaded with the following command:

```
# service auditd reload
```

4.1.1 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.

4.1.1.1 Ensure auditd is installed (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

`auditd` is the user space 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 audit-libs
```

Remediation:

Run the following command to Install `auditd`:

```
# dnf install audit audit-libs
```

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	6.2 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.			

4.1.1.2 Ensure auditd service is enabled (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

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
```

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	8.5 <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.		●	●

4.1.1.3 Enable auditing for processes that start prior to auditd (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

ziPL is the default Red Hat Enterprise Linux 8 boot loader for IBM Z systems. Ensure that all processes can be audited, even those which start prior to the audit daemon (`auditd`).

Make sure `/etc/zipl.conf` doesn't contain `image = setting`, as Red Hat Enterprise Linux 8 adheres to Boot Loader Specification (BLS).

Run the `zipl` command so that `/boot/bootmap` is updated.

Rationale:

Audit events need to be captured on processes that start up prior to `auditd`, so that potential malicious activity cannot go undetected.

Impact:

Each process on the system carries an `auditable` flag which indicates whether its activities can be audited. Although the Linux Auditing System `auditd` takes care of enabling this flag for all processes which launch after `auditd` launches, adding the kernel argument (adding `audit=1` to `/etc/kernel/cmdline`) ensures that the flag is set for every process during boot.

Audit:

Check that all boot entries in the zipl boot loader

```
/boot/loader/entries/*.conf
```

have `audit=1` included in their options.

Remediation:

To ensure that new kernels and boot entries continue to enable audit, add the entry `audit=1` to file

```
/etc/kernel/cmdline
```

Default Value:

`audit=1`

Additional Information:

This recommendation is designed around the grub2 boot loader, if another boot loader is in use in your environment enact equivalent settings.

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.		●	●

4.1.1.4 Extend the audit backlog limit for the audit daemon (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

To improve the kernel capacity to queue all log events, even those which start prior to the audit daemon, check that all boot entries in file

```
/boot/loader/entries/*.conf
```

have the entry `audit_backlog_limit=8192` included in their options.

Make sure that file

```
/etc/zipl.conf
```

does not contain the entry `image = setting`, because Red Hat Enterprise Linux 8 adheres to the **Boot Loader Specification (BLS)**.

Additionally, run the `zipl` command so that

```
/boot/bootmap
```

is updated.

Rationale:

If during boot, `audit=1`, the backlog holds 64 records.

If more than 64 records are created during boot, `auditd` records are lost and potential malicious activity may remain undetected.

Impact:

The `audit_backlog_limit` sets the queue length for audit events awaiting transfer to the audit daemon. Until the audit daemon is up and running, all log messages are stored in this queue. If the queue overruns during boot process, the action defined by the audit failure flag is taken.

Audit:

All boot entries in file

```
/boot/loader/entries/*.conf
```

have the queue length set to: `audit_backlog_limit=8192`

Remediation:

Check that all boot entries in file

```
/boot/loader/entries/*.conf
```

have the entry `audit_backlog_limit=8192`
included in their options.
Make sure that file

```
/etc/zipl.conf
```

does not contain the entry `image = setting,`
because Red Hat Enterprise Linux 8 adheres to the
Boot Loader Specification (BLS).
Additionally, run the `zipl` command so that

```
/boot/bootmap
```

is updated.
To ensure that new kernels and boot entries
continue to extend the audit log events queue,
add the `audit_backlog_limit=8192` option to file

```
/etc/kernel/cmdline
```

Default Value:

30

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.		●	●

4.1.2 Configure Data Retention

When auditing, it is important to carefully configure the storage requirements for audit logs. By default, `auditd` maximizes the log files at 5MB and retains only 4 copies of them. Older versions are 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.

4.1.2.1 Ensure audit log storage size is configured (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

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 max_log_file /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>
```

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.		●	●

4.1.2.2 Ensure audit logs are not automatically deleted (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

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
```

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.		●	●

4.1.2.3 Ensure system is disabled when audit logs are full (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

The `auditd` daemon can be configured to halt the system when the audit logs are full.

Rationale:

In high security contexts, the risk of detecting unauthorized access or nonrepudiation exceeds the benefit of the system's availability.

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
# grep admin_space_left_action /etc/audit/auditd.conf

admin_space_left_action = halt
```

Remediation:

Set the following parameters in `/etc/audit/auditd.conf`:

```
space_left_action = email
action_mail_acct = root
admin_space_left_action = halt
```

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.			

4.1.3 Ensure changes to system administration scope (sudoers) is collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

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` will be written to when the file or its attributes have changed. The audit records will be tagged with the identifier "scope."

Rationale:

Changes in the `/etc/sudoers` file can indicate that an unauthorized change has been made to scope of system administrator activity.

Audit:

Run the following commands:

```
# grep scope /etc/audit/rules.d/*.rules
# auditctl -l | grep scope
```

Verify output of both 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`

Example: `vi /etc/audit/rules.d/scope.rules`

and add the following lines:

```
-w /etc/sudoers -p wa -k scope
-w /etc/sudoers.d/ -p wa -k scope
```

Additional Information:

Reloading the `auditd` config to set active settings may require a system reboot.

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.			

4.1.4 Ensure login and logout events are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Monitor login and logout events. The parameters below track changes to files associated with login/logout events. The file `/var/log/faillog` tracks failed events from login. The file `/var/log/lastlog` maintain records of the last time a user successfully logged in.

Rationale:

Monitoring login/logout events could provide a system administrator with information associated with brute force attacks against user logins.

Audit:

Run the following commands:

```
# grep logins /etc/audit/rules.d/*.rules
# auditctl -l | grep logins
```

Verify output of both includes:

```
-w /var/log/faillog -p wa -k logins
-w /var/log/lastlog -p wa -k logins
```

Remediation:

Edit or create a file in the `/etc/audit/rules.d/` directory ending in `.rules`

Example: `vi /etc/audit/rules.d/audit.rules`

and add the following lines:

```
-w /var/log/faillog -p wa -k logins
-w /var/log/lastlog -p wa -k logins
```

Additional Information:

Reloading the `auditd` config to set active settings may require a system reboot.

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.			

4.1.5 Ensure session initiation information is collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Monitor session initiation events. The parameters in this section track changes to the files associated with session events. The file `/var/run/utmp` tracks all currently logged in users. All audit records will be tagged with the identifier "session." The `/var/log/wtmp` file tracks logins, logouts, shutdown, and reboot events. The file `/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 "logins."

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:

Run the following commands:

```
# grep -E '(session|logins)' /etc/audit/rules.d/*.rules
# auditctl -l | grep -E '(session|logins)'
```

Verify output of both includes:

```
-w /var/run/utmp -p wa -k session
-w /var/log/wtmp -p wa -k logins
-w /var/log/btmp -p wa -k logins
```

Remediation:

Edit or create a file in the `/etc/audit/rules.d/` directory ending in `.rules`

Example: `vi /etc/audit/rules.d/logins.rules`

and add the following lines:

```
-w /var/run/utmp -p wa -k session
-w /var/log/wtmp -p wa -k logins
-w /var/log/btmp -p wa -k logins
```

Additional Information:

The `last` command can be used to read `/var/log/wtmp` (`last` with no parameters) and `/var/run/utmp` (`last -f /var/run/utmp`)

Reloading the `auditd` config to set active settings may require a system reboot.

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.			

4.1.6 Ensure events that modify date and time information are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

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) or `clock_settime` (allows for the setting of several internal clocks and timers) system calls have been executed and always write an audit record to the `/var/log/audit.log` file upon exit, tagging the records with the identifier "time-change"

Rationale:

Unexpected changes in system date and/or time could be a sign of malicious activity on the system.

Audit:

Run the following commands:

```
# grep time-change /etc/audit/rules.d/*.rules
# auditctl -l | grep time-change
```

Verify output of both matches:

```
-a always,exit -F arch=b64 -S adjtimex -S settimeofday -k time-change
-a always,exit -F arch=b32 -S adjtimex -S settimeofday -S stime -k time-change
-a always,exit -F arch=b64 -S clock_settime -k time-change
-a always,exit -F arch=b32 -S clock_settime -k time-change
-w /etc/localtime -p wa -k time-change
```

Remediation:

Edit or create a file in the `/etc/audit/rules.d/` directory ending in `.rules`

Example: `vi /etc/audit/rules.d/time-change.rules`

and add the following lines:

```
-a always,exit -F arch=b64 -S adjtimex -S settimeofday -k time-change
-a always,exit -F arch=b32 -S adjtimex -S settimeofday -S stime -k time-change
-a always,exit -F arch=b64 -S clock_settime -k time-change
-a always,exit -F arch=b32 -S clock_settime -k time-change
-w /etc/localtime -p wa -k time-change
```

Additional Information:

Reloading the `auditd` config to set active settings may require a system reboot.

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.			

4.1.7 Ensure events that modify the system's Mandatory Access Controls are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Monitor SELinux/AppArmor 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 or /etc/apparmor and /etc/apparmor.d directories.

Rationale:

Changes to files in these 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 systems using SELinux run the following commands:

```
# grep MAC-policy /etc/audit/rules.d/*.rules
# auditctl -l | grep MAC-policy
```

Verify output of both matches:

```
-w /etc/selinux/ -p wa -k MAC-policy
-w /usr/share/selinux/ -p wa -k MAC-policy
```

Remediation:

On systems using SELinux Edit or create a file in the /etc/audit/rules.d/ directory ending in .rules

Example: vi /etc/audit/rules.d/MAC-policy.rules
and add the following lines:

```
-w /etc/selinux/ -p wa -k MAC-policy
-w /usr/share/selinux/ -p wa -k MAC-policy
```

Additional Information:

Reloading the auditd config to set active settings may require a system reboot.

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.			

4.1.8 Ensure events that modify the system's network environment are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Record changes to network environment files or system calls. The below parameters monitor the `sethostname` (set the systems host name) or `setdomainname` (set the systems domainname) system calls, and write an audit event on system call exit. The other parameters monitor the `/etc/issue` and `/etc/issue.net` files (messages displayed pre-login), `/etc/hosts` (file containing host names and associated IP addresses) and `/etc/sysconfig/network` (directory containing network interface scripts and configurations) files.

Rationale:

Monitoring `sethostname` and `setdomainname` will identify potential unauthorized changes to host and domainname 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 in the file 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 will be tagged with the identifier "system-locale."

Audit:

Run the following commands:

```
# grep system-locale /etc/audit/rules.d/*.rules
# auditctl -l | grep system-locale
```

Verify output of both matches:

```
-a always,exit -F arch=b64 -S sethostname -S setdomainname -k system-locale
-a always,exit -F arch=b32 -S sethostname -S 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
```

Remediation:

Edit or create a file in the `/etc/audit/rules.d/` directory ending in `.rules`

Example: `vi /etc/audit/rules.d/system-locale.rules`

and add the following lines:

```
-a always,exit -F arch=b64 -S sethostname -S setdomainname -k system-locale
-a always,exit -F arch=b32 -S sethostname -S 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
```

Additional Information:

Reloading the `auditd` config to set active settings may require a system reboot.

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.			

4.1.9 Ensure discretionary access control permission modification events are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

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 `chmod`, `fchmod` and `fchmodat` system calls affect the permissions associated with a file. The `chown`, `fchown`, `fchownat` and `lchown` system calls affect owner and group attributes on a file. The `setxattr`, `lsetxattr`, `fsetxattr` (set extended file attributes) and `removexattr`, `lremovexattr`, `fremovexattr` (remove extended file attributes) control extended file attributes. In all cases, an audit record will only be written for non-system user ids (`audit >= 1000`) and will ignore Daemon events (`audit = 4294967295`). All audit records will be tagged with the identifier "perm_mod."

Note: Systems may have been customized to change the default `UID_MIN`. To confirm the `UID_MIN` for your system, run the following command:

```
awk '/^\s*UID_MIN/{print $2}' /etc/login.defs
```

If your systems' `UID_MIN` is not `1000`, replace `audit>=1000` with `audit>=<UID_MIN for your system>` in the Audit and Remediation procedures.

Rationale:

Monitoring for changes in file attributes could alert a system administrator to activity that could indicate intruder activity or policy violation.

Audit:

Run the following commands:

```
# grep perm_mod /etc/audit/rules.d/*.rules
```

Verify output matches:

```
-a always,exit -F arch=b64 -S chmod -S fchmod -S fchmodat -F auid>=1000 -F auid!=4294967295 -k perm_mod
-a always,exit -F arch=b32 -S chmod -S fchmod -S fchmodat -F auid>=1000 -F auid!=4294967295 -k perm_mod
-a always,exit -F arch=b64 -S chown -S fchown -S fchownat -S lchown -F auid>=1000 -F auid!=4294967295 -k perm_mod
-a always,exit -F arch=b32 -S chown -S fchown -S fchownat -S lchown -F auid>=1000 -F auid!=4294967295 -k perm_mod
-a always,exit -F arch=b64 -S setxattr -S lsetxattr -S fsetxattr -S removexattr -S lremovexattr -S fremovexattr -F auid>=1000 -F auid!=4294967295 -k perm_mod
-a always,exit -F arch=b32 -S setxattr -S lsetxattr -S fsetxattr -S removexattr -S lremovexattr -S fremovexattr -F auid>=1000 -F auid!=4294967295 -k perm_mod
# auditctl -l | grep perm_mod
```

Verify output matches:

```
-a always,exit -F arch=b64 -S chmod -S fchmod -S fchmodat -F auid>=1000 -F auid!=-1 -k perm_mod
-a always,exit -F arch=b32 -S chmod -S fchmod -S fchmodat -F auid>=1000 -F auid!=-1 -k perm_mod
-a always,exit -F arch=b64 -S chown -S fchown -S fchownat -S lchown -F auid>=1000 -F auid!=-1 -k perm_mod
-a always,exit -F arch=b32 -S chown -S fchown -S fchownat -S lchown -F auid>=1000 -F auid!=-1 -k perm_mod
-a always,exit -F arch=b64 -S setxattr -S lsetxattr -S fsetxattr -S removexattr -S lremovexattr -S fremovexattr -F auid>=1000 -F auid!=-1 -k perm_mod
-a always,exit -F arch=b32 -S setxattr -S lsetxattr -S fsetxattr -S removexattr -S lremovexattr -S fremovexattr -F auid>=1000 -F auid!=-1 -k perm_mod
```

Remediation:

Edit or create a file in the `/etc/audit/rules.d/` directory ending in `.rules`

Example: `vi /etc/audit/rules.d/perm_mod.rules`

and add the following lines:

```
-a always,exit -F arch=b64 -S chmod -S fchmod -S fchmodat -F auid>=1000 -F auid!=4294967295 -k perm_mod
-a always,exit -F arch=b32 -S chmod -S fchmod -S fchmodat -F auid>=1000 -F auid!=4294967295 -k perm_mod
-a always,exit -F arch=b64 -S chown -S fchown -S fchownat -S lchown -F auid>=1000 -F auid!=4294967295 -k perm_mod
-a always,exit -F arch=b32 -S chown -S fchown -S fchownat -S lchown -F auid>=1000 -F auid!=4294967295 -k perm_mod
-a always,exit -F arch=b64 -S setxattr -S lsetxattr -S fsetxattr -S removexattr -S lremovexattr -S fremovexattr -F auid>=1000 -F auid!=4294967295 -k perm_mod
-a always,exit -F arch=b32 -S setxattr -S lsetxattr -S fsetxattr -S removexattr -S lremovexattr -S fremovexattr -F auid>=1000 -F auid!=4294967295 -k perm_mod
```

Additional Information:

Reloading the `auditd` config to set active settings may require a system reboot.

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.			

4.1.10 Ensure unsuccessful unauthorized file access attempts are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Monitor for unsuccessful attempts to access files. The parameters below are associated with system calls that control creation (`creat`), opening (`open` , `openat`) and truncation (`truncate` , `ftruncate`) of files. An audit log record will only be written if the user is a non-privileged user (`audit >= 1000`), is not a Daemon event (`audit=4294967295`) and if the system call returned `EACCES` (permission denied to the file) or `EPERM` (some other permanent error associated with the specific system call). All audit records will be tagged with the identifier "access."

Note: Systems may have been customized to change the default `UID_MIN`. To confirm the `UID_MIN` for your system, run the following command:

```
awk '/^\s*UID_MIN/{print $2}' /etc/login.defs
```

If your systems' `UID_MIN` is not `1000`, replace `audit>=1000` with `audit>=<UID_MIN for your system>` in the Audit and Remediation procedures.

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:

Run the following commands:

```
# grep access /etc/audit/rules.d/*.rules
```

Verify output matches:

```
-a always,exit -F arch=b64 -S creat -S open -S openat -S truncate -S
ftruncate -F exit=-EACCES -F auid>=1000 -F auid!=4294967295 -k access
-a always,exit -F arch=b32 -S creat -S open -S openat -S truncate -S
ftruncate -F exit=-EACCES -F auid>=1000 -F auid!=4294967295 -k access
-a always,exit -F arch=b64 -S creat -S open -S openat -S truncate -S
ftruncate -F exit=-EPERM -F auid>=1000 -F auid!=4294967295 -k access
-a always,exit -F arch=b32 -S creat -S open -S openat -S truncate -S
ftruncate -F exit=-EPERM -F auid>=1000 -F auid!=4294967295 -k access
# auditctl -l | grep access
```

Verify output matches:

```
-a always,exit -F arch=b64 -S creat -S open -S openat -S truncate -S
ftruncate -F exit=-EACCES -F auid>=1000 -F auid!=-1 -k access
-a always,exit -F arch=b32 -S creat -S open -S openat -S truncate -S
ftruncate -F exit=-EACCES -F auid>=1000 -F auid!=-1 -k access
-a always,exit -F arch=b64 -S creat -S open -S openat -S truncate -S
ftruncate -F exit=-EPERM -F auid>=1000 -F auid!=-1 -k access
-a always,exit -F arch=b32 -S creat -S open -S openat -S truncate -S
ftruncate -F exit=-EPERM -F auid>=1000 -F auid!=-1 -k access
```

Remediation:

Edit or create a file in the `/etc/audit/rules.d/` directory ending in `.rules`

Example: `vi /etc/audit/rules.d/access.rules`

and add the following lines:

```
-a always,exit -F arch=b64 -S creat -S open -S openat -S truncate -S
ftruncate -F exit=-EACCES -F auid>=1000 -F auid!=4294967295 -k access
-a always,exit -F arch=b32 -S creat -S open -S openat -S truncate -S
ftruncate -F exit=-EACCES -F auid>=1000 -F auid!=4294967295 -k access
-a always,exit -F arch=b64 -S creat -S open -S openat -S truncate -S
ftruncate -F exit=-EPERM -F auid>=1000 -F auid!=4294967295 -k access
-a always,exit -F arch=b32 -S creat -S open -S openat -S truncate -S
ftruncate -F exit=-EPERM -F auid>=1000 -F auid!=4294967295 -k access
```

Additional Information:

Reloading the `auditd` config to set active settings may require a system reboot.

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	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).			

4.1.11 Ensure events that modify user/group information are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Record events affecting the `group`, `passwd` (user IDs), `shadow` and `gshadow` (passwords) or `/etc/security/opasswd` (old passwords, based on remember parameter in the PAM configuration) files. 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:

Run the following commands:

```
# grep identity /etc/audit/rules.d/*.rules
# auditctl -l | grep identity
```

Verify output of both 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`

Example: `vi /etc/audit/rules.d/identity.rules`

and add the following lines:

```
-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
```

Additional Information:

Reloading the `auditd` config to set active settings may require a system reboot.

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.			

4.1.12 Ensure successful file system mounts are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

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.

Note: Systems may have been customized to change the default `UID_MIN`. To confirm the `UID_MIN` for your system, run the following command:

```
awk '/^\s*UID_MIN/{print $2}' /etc/login.defs
```

If your systems' `UID_MIN` is not `1000`, replace `audit>=1000` with `audit>=<UID_MIN for your system>` in the Audit and Remediation procedures.

Audit:

Run the following commands:

```
# grep mounts /etc/audit/rules.d/*.rules
```

Verify output matches:

```
-a always,exit -F arch=b64 -S mount -F auid>=1000 -F auid!=4294967295 -k mounts
-a always,exit -F arch=b32 -S mount -F auid>=1000 -F auid!=4294967295 -k mounts
# auditctl -l | grep mounts
```

Verify output matches:

```
-a always,exit -F arch=b64 -S mount -F auid>=1000 -F auid!=-1 -k mounts
-a always,exit -F arch=b32 -S mount -F auid>=1000 -F auid!=-1 -k mounts
```

Remediation:

Edit or create a file in the `/etc/audit/rules.d/` directory ending in `.rules`

Example: `vi /etc/audit/rules.d/mounts.rules`

and add the following lines:

```
-a always,exit -F arch=b64 -S mount -F auid>=1000 -F auid!=4294967295 -k mounts
-a always,exit -F arch=b32 -S mount -F auid>=1000 -F auid!=4294967295 -k mounts
```

Additional Information:

This tracks successful and unsuccessful mount commands. File system mounts do not have to come from external media and this action still does not verify write (e.g. CD ROMS).

Reloading the `auditd` config to set active settings may require a system reboot.

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	6.2 Activate audit logging Ensure that local logging has been enabled on all systems and networking devices.	●	●	●

4.1.13 Ensure use of privileged commands is collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

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.

Audit:

Run the following command replacing <partition> with a list of partitions where programs can be executed from on your system:

```
# find <partition> -xdev \( -perm -4000 -o -perm -2000 \) -type f | awk
'{print "-a always,exit -F path=" $1 " -F perm=x -F auid>=" $(awk
'/^\s*UID_MIN/{print $2}' /etc/login.defs) "' -F auid!=4294967295 -k
privileged" }'
```

Verify all resulting lines are in a .rules file in /etc/audit/rules.d/ and the output of auditctl -l.

Note: The .rules file output will be auid!=-1 not auid!=4294967295

Remediation:

To remediate this issue, the system administrator will have to execute a find command to locate all the privileged programs and then add an audit line for each one of them. The audit parameters associated with this are as follows:

-F path=" \$1 " - will populate each file name found through the find command and processed by awk. -F perm=x - will write an audit record if the file is executed. -F auid>=1000 - will write a record if the user executing the command is not a privileged user. -F auid!= 4294967295 - will ignore Daemon events

All audit records should be tagged with the identifier "privileged".

Run the following command replacing with a list of partitions where programs can be executed from on your system:

```
# find <partition> -xdev \( -perm -4000 -o -perm -2000 \) -type f | awk
'{print "-a always,exit -F path=" $1 " -F perm=x -F auid>="$(awk
'/^\s*UID_MIN/{print $2}' /etc/login.defs) "' -F auid!=4294967295 -k
privileged" }'
```

Edit or create a file in the /etc/audit/rules.d/ directory ending in .rules and add all resulting lines to the file.

Example:

```
# find <partition> -xdev \( -perm -4000 -o -perm -2000 \) -type f | awk
'{print "-a always,exit -F path=" $1 " -F perm=x -F auid>="$(awk
'/^\s*UID_MIN/{print $2}' /etc/login.defs) "' -F auid!=4294967295 -k
privileged" }' >> /etc/audit/rules.d/privileged.rules
```

Additional Information:

Reloading the auditd config to set active settings may require a system reboot.

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	6.2 Activate audit logging Ensure that local logging has been enabled on all systems and networking devices.	●	●	●

4.1.14 Ensure file deletion events by users are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

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 the `unlink` (remove a file), `unlinkat` (remove a file attribute), `rename` (rename a file) and `renameat` (rename a file attribute) system calls and tags them with the identifier "delete".

Note: Systems may have been customized to change the default `UID_MIN`. To confirm the `UID_MIN` for your system, run the following command:

```
awk '/^\s*UID_MIN/{print $2}' /etc/login.defs
```

If your systems' `UID_MIN` is not 1000, replace `audit>=1000` with `audit>=<UID_MIN for your system>` in the Audit and Remediation procedures.

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:

Run the following commands:

```
# grep delete /etc/audit/rules.d/*.rules
```

Verify output matches:

```
-a always,exit -F arch=b64 -S unlink -S unlinkat -S rename -S renameat -F
auid>=1000 -F auid!=4294967295 -k delete
-a always,exit -F arch=b32 -S unlink -S unlinkat -S rename -S renameat -F
auid>=1000 -F auid!=4294967295 -k delete
# auditctl -l | grep delete
```

Verify output matches:

```
-a always,exit -F arch=b64 -S unlink -S unlinkat -S rename -S renameat -F
auid>=1000 -F auid!=-1 -k delete
-a always,exit -F arch=b32 -S unlink -S unlinkat -S rename -S renameat -F
auid>=1000 -F auid!=-1 -k delete
```

Remediation:

Edit or create a file in the `/etc/audit/rules.d/` directory ending in `.rules`

Example: `vi /etc/audit/rules.d/delete.rules`

and add the following lines:

```
-a always,exit -F arch=b64 -S unlink -S unlinkat -S rename -S renameat -F
auid>=1000 -F auid!=4294967295 -k delete
-a always,exit -F arch=b32 -S unlink -S unlinkat -S rename -S renameat -F
auid>=1000 -F auid!=4294967295 -k delete
```

Additional Information:

At a minimum, configure the audit system to collect file deletion events for all users and root.

Reloading the auditd config to set active settings may require a system reboot.

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.	●	●	●

4.1.15 Ensure kernel module loading and unloading is collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Monitor the loading and unloading of kernel modules. The programs `insmod` (install a kernel module), `rmmod` (remove a kernel module), and `modprobe` (a more sophisticated program to load and unload modules, as well as some other features) control loading and unloading of modules. The `init_module` (load a module) and `delete_module` (delete a module) system calls control loading and unloading of 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 `insmod`, `rmmod` and `modprobe` could provide system administrators with evidence that an unauthorized user loaded or unloaded a kernel module, possibly compromising the security of the system. Monitoring of the `init_module` and `delete_module` system calls would reflect an unauthorized user attempting to use a different program to load and unload modules.

Audit:

Run the following commands:

```
# grep modules /etc/audit/rules.d/*.rules
# auditctl -l | grep modules
```

Verify output of both matches:

```
-w /sbin/insmod -p x -k modules
-w /sbin/rmmod -p x -k modules
-w /sbin/modprobe -p x -k modules
-a always,exit -F arch=b64 -S init_module -S delete_module -k modules
```

Remediation:

Edit or create a file in the `/etc/audit/rules.d/` directory ending in `.rules`

Example: `vi /etc/audit/rules.d/modules.rules`

and add the following lines:

```
-w /sbin/insmod -p x -k modules
-w /sbin/rmmod -p x -k modules
-w /sbin/modprobe -p x -k modules
-a always,exit -F arch=b64 -S init_module -S delete_module -k modules
```

Additional Information:

Reloading the `auditd` config to set active settings may require a system reboot.

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	6.2 Activate audit logging Ensure that local logging has been enabled on all systems and networking devices.			

4.1.16 Ensure system administrator actions (sudolog) are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

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:

Run the following commands:

```
# grep -E "^\\s*-w\\s+$(grep -r logfile /etc/sudoers* | sed -e  
's/.*/logfile=//;s/,? .*//')\\s+-p\\s+wa\\s+-k\\s+actions"  
/etc/audit/rules.d/*.rules  
  
# auditctl -l | grep actions
```

Verify output of both matches the output of the following command, and the the output includes a file path

```
echo "-w $(grep -r logfile /etc/sudoers* | sed -e 's/.*/logfile=//;s/,? .*//')  
-p wa -k actions"
```

Example Output

```
-w /var/log/sudo.log -p wa -k actions
```

Remediation:

Edit or create a file in the `/etc/audit/rules.d/` directory ending in `.rules` and add the following line:

```
-w <Path to sudo logfile> -p wa -k actions
```

Example: `vi /etc/audit/rules.d/audit.rules`
and add the following line:

```
-w /var/log/sudo.log -p wa -k actions
```

Additional Information:

The system must be configured with `su` disabled (See Item 5.6 Ensure access to the `su` command is restricted) to force all command execution through `sudo`. This will not be effective on the console, as administrators can log in as root.

Reloading the `auditd` config to set active settings may require a system reboot.

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.			

4.1.17 Ensure the audit configuration is immutable (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

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.

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 "^s*[^#]" /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

Additional Information:

This setting will ensure reloading the `auditd` config to set active settings requires a system reboot.

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.			

4.2 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 and ease log analysis.

4.2.1 Configure rsyslog

The `rsyslog` software is recommended as a replacement for the `syslogd` daemon and provides improvements over `syslogd`, 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. **Note:** This section only applies if `rsyslog` is installed on the system.

4.2.1.1 Ensure rsyslog is installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `rsyslog` software is a recommended replacement to the original `syslogd` daemon which provide improvements over `syslogd`, 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.

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:

Verify rsyslog is installed.
Run the following command

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

Remediation:

Run the following command to install rsyslog:

```
# dnf install rsyslog
```

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.			

4.2.1.2 Ensure rsyslog Service is enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Once the `rsyslog` package is installed it needs to be activated.

Rationale:

If the `rsyslog` service is not activated the system may default to the `syslogd` service or lack logging instead.

Audit:

Run the following command to verify `rsyslog` is enabled:

```
# systemctl is-enabled rsyslog
enabled
```

Verify result is "enabled".

Remediation:

Run the following command to enable `rsyslog`:

```
# systemctl --now enable rsyslog
```

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.			

4.2.1.3 Ensure rsyslog default file permissions configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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.

Audit:

Run the following command and verify that `$FileCreateMode` is 0640 or more restrictive:

```
# grep ^\ $FileCreateMode /etc/rsyslog.conf /etc/rsyslog.d/*.conf
```

Remediation:

Edit the `/etc/rsyslog.conf` and `/etc/rsyslog.d/*.conf` files and set `$FileCreateMode` to 0640 or more restrictive:

```
$FileCreateMode 0640
```

References:

1. See the `rsyslog.conf(5)` man page for more information.

Additional Information:

You should also ensure this is not overridden with less restrictive settings in any `/etc/rsyslog.d/*` conf file.

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	5.1 <u>Establish Secure Configurations</u> Maintain documented, standard security configuration standards for all authorized operating systems and software.			

4.2.1.4 Ensure logging is configured (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 the `/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:

```
# 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:

```
*.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
news.crit                              -/var/log/news/news.crit
news.err                               -/var/log/news/news.err
news.notice                           -/var/log/news/news.notice
*.=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.

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.			

4.2.1.5 Ensure rsyslog is configured to send logs to a remote log host (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `rsyslog` utility supports the ability to send logs it gathers to a remote log host running `syslogd(8)` or to receive messages from remote hosts, reducing administrative overhead.

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):

```
# grep "^*.*[^\I][^\I]*@" /etc/rsyslog.conf /etc/rsyslog.d/*.conf
*.* @@loghost.example.com
```

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).

```
*.* @@loghost.example.com
```

Run the following command to reload the `rsyslogd` configuration:

```
# systemctl restart rsyslog
```

References:

1. See the `rsyslog.conf(5)` man page for more information.

Additional Information:

The double "at" sign (`@@`) directs `rsyslog` to use TCP to send log messages to the server, which is a more reliable transport mechanism than the default UDP protocol.

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	6.6 <u>Deploy SIEM or Log Analytic tool</u> Deploy Security Information and Event Management (SIEM) or log analytic tool for log correlation and analysis.		●	●
v7	6.8 <u>Regularly Tune SIEM</u> On a regular basis, tune your SIEM system to better identify actionable events and decrease event noise.			●

4.2.1.6 Ensure remote rsyslog messages are only accepted on designated log hosts. (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

By default, `rsyslog` does not listen for log messages coming in from remote systems. The `ModLoad` tells `rsyslog` to load the `imtcp.so` module so it can listen over a network via TCP. The `InputTCPServerRun` option instructs `rsyslogd` to listen on the specified TCP port.

Rationale:

The guidance in the section ensures that remote log hosts are configured to only accept `rsyslog` data from hosts within the specified domain and that those systems that are not designed to be log hosts do not accept any remote `rsyslog` messages. This provides protection from spoofed log data and ensures that system administrators are reviewing reasonably complete syslog data in a central location.

Audit:

Run the following commands and verify the resulting lines are uncommented on designated log hosts and commented or removed on all others:

```
# grep '$ModLoad imtcp' /etc/rsyslog.conf /etc/rsyslog.d/*.conf

$ModLoad imtcp
# grep '$InputTCPServerRun' /etc/rsyslog.conf /etc/rsyslog.d/*.conf

$InputTCPServerRun 514
```

Remediation:

For hosts that are designated as log hosts, edit the `/etc/rsyslog.conf` file and un-comment or add the following lines:

```
$ModLoad imtcp
$InputTCPServerRun 514
```

For hosts that are not designated as log hosts, edit the `/etc/rsyslog.conf` file and comment or remove the following lines:

```
# $ModLoad imtcp
# $InputTCPServerRun 514
```

Run the following command to reload the `rsyslogd` configuration:

```
# systemctl restart rsyslog
```

References:

1. See the `rsyslog(8)` man page for more information.

Additional Information:

The `$ModLoad imtcp` line can have the `.so` extension added to the end of the module, or use the full path to the module.

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.			

4.2.2 Configure journald

systemd-journald is a system service that collects and stores logging data. It creates and maintains structured, indexed journals based on logging information that is received from a variety of sources: Kernel log messages, via kmsg

Any changes made to the systemd-journald configuration will require a re-start of systemd-journald

4.2.2.1 Ensure journald is configured to send logs to rsyslog (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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:

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 `/etc/systemd/journald.conf` and verify that logs are forwarded to syslog

```
# grep -e ^\s*ForwardToSyslog /etc/systemd/journald.conf
ForwardToSyslog=yes
```

Remediation:

Edit the `/etc/systemd/journald.conf` file and add the following line:

```
ForwardToSyslog=yes
```

References:

1. <https://github.com/konstruktoid/hardening/blob/master/systemd.adoc#etcsystemdjournaldconf>

Additional Information:

This recommendation assumes that recommendation 4.2.1.5, "Ensure rsyslog is configured to send logs to a remote log host" has been implemented.

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-journald.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-journald.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/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	6.5 <u>Central Log Management</u> Ensure that appropriate logs are being aggregated to a central log management system for analysis and review.			

4.2.2.2 Ensure journald is configured to compress large log files (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 -e ^\s*Compress /etc/systemd/journald.conf
Compress=yes
```

Remediation:

Edit the `/etc/systemd/journald.conf` file and add the following line:

```
Compress=yes
```

References:

1. <https://github.com/konstruktoid/hardening/blob/master/systemd.adoc#etcsystemdjournaldconf>

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	6.4 <u>Ensure adequate storage for logs</u> Ensure that all systems that store logs have adequate storage space for the logs generated.			

4.2.2.3 Ensure journald is configured to write logfiles to persistent disk (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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.

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 -e ^\s*Storage /etc/systemd/journald.conf
Storage=persistent
```

Remediation:

Edit the `/etc/systemd/journald.conf` file and add the following line:

```
Storage=persistent
```

References:

1. <https://github.com/konstruktoid/hardening/blob/master/systemd.adoc#etcsystemdjournaldconf>

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.			

4.2.3 *Ensure permissions on all logfiles are configured (Automated)*

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Log files stored in /var/log/ contain logged information from many services on the system, or on log hosts others as well.

Rationale:

It is important to ensure that log files have the correct permissions to ensure that sensitive data is archived and protected.

Audit:

Run the following command and verify that other has no permissions on any files and group does not have write or execute permissions on any files:

```
# find /var/log -type f -perm /037 -ls -o -type d -perm /026 -ls
```

No output should be returned

Remediation:

Run the following commands to set permissions on all existing log files:

```
find /var/log -type f -exec chmod g-wx,o-rwx "{}" + -o -type d -exec chmod g-w,o-rwx "{}" +
```

Additional Information:

You may also need to change the configuration for your logging software or services for any logs that had incorrect permissions.

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.			

4.2.4 Ensure permissions on operating system resources (OSR) (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

You need to protect the operating system resources - OSR files - in directory

```
/var/log/secure.
```

Set the file permissions as follows:

```
rwX r-- ---
```

(or even more restrictive)

Notes:

```
/var/log/secure
```

is required on all Linux distributions except Debian.

```
/var/log/auth.log
```

is required on Debian Linux.

Rationale:

Impact:

You need to ensure that the

```
/var/log/secure
```

resources have the correct permission.

Audit:

File permissions must be set at least as shown:

```
rwX r-- ---
```

(or even more restrictive).

Run the following command to verify the file permissions:

```
ls -lt /var/log/secure
```

and check if file permissions are sufficient.

The output may be similar to the following:

```
-rw-----. 1 root root 1239 Jul 30 12:01 /var/log/secure
```

Remediation:

Run the following command to set up correct permission:

```
chmod g-wx,o-wrx /var/log/secure
```

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.			

4.3 Ensure logrotate is configured (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 very 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.

Additional Information:

If no `maxage` setting is set for logrotate a situation can occur where logrotate is interrupted and fails to delete rotated logfiles. 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 logfile is removed but standard rotation settings are not overridden.

CIS Controls:

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

4.4 Ensure logging of login success or failure (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Your system should provide the logging of login success or failure by writing according log messages to specified files. Read about the requirements for systems that use `rsyslog` to ensure that needed login success or failure messages are captured.

Check whether the file `/etc/rsyslog.conf` contains the following entries:

```
$ActionFileDefaultTemplate RSYSLOG_TraditionalFileFormat
.info;mail.none;authpriv.none;cron.none /var/log/messages
authpriv. /var/log/secure
```

Or, if you do not want to use the global variable `$ActionFileDefaultTemplate`, your `/etc/rsyslog.conf` should look similar to the following:

```
.info;mail.none;authpriv.none;cron.none /var/log/messages;
RSYSLOG_TraditionalFileFormat
authpriv. /var/log/secure;RSYSLOG_TraditionalFileFormat
```

The use of the '-' sign in `-/var/log/...`, which indicates that buffered writes are allowed, is optional in any Linux `syslog` configuration.

`rsyslog` supports a `$IncludeConfig` directive to add settings using files usually located in the `/etc/rsyslog.d/` directory, where the required settings can be included using this directive.

More extensive logging is permissible, for example, `cron.none` can be removed from the facility section of the control stanza for `/var/log/messages`.

If your account selects to use `syslog-ng` instead of `rsyslog`, ensure that the **agreed values** align with your account's expectations for `syslog-ng`.

Rationale:

Impact:

Important system status messages are captured, if you use the described settings.

Audit:

To verify that the needed records are captured via `rsyslog`, ensure that the following two checking routines finish successfully:

- The following three `grep` checks must pass for file `/etc/rsyslog.conf`

1. Enter the following command:

```
grep '^[^#]$ActionFileDefaultTemplate\sRSYSLOG_TraditionalFileFormat' $file
```

This command must return

```
$ActionFileDefaultTemplate RSYSLOG_TraditionalFileFormat
```

2. Enter the following command:

```
grep '^[^#]*.info;mail.none;authpriv.none;cron.none\s[-]* /var/log/messages' $file
```

This command must return must return

```
*.info;mail.none;authpriv.none;cron.none  
-/var/log/messages
```

The '-' sign in front of the file path is optional.

3. Enter the following command:

```
egrep '^[^#]authpriv.*\s[-]* /var/log/secure' $file
```

This command must return

```
authpriv.* -/var/log/secure
```

The '-' sign in front of the file path is optional.

- The following two checks must also return the desired results:

1. Enter the following command:

```
grep '^[^#]*.info;mail.none;authpriv.none;cron.none\s[-]* /var/log/messages' $file
```

This command must return


```
*.info;mail.none;authpriv.none;cron.none
-/var/log/messages;RSYSLOG_TraditionalFileFormat
```

The '-' sign in front of the file path is optional.

2. Enter the following command:

```
egrep '^[^#]authpriv.*\s[-]*/var/log/secure' $file
```

This command must return

```
authpriv.* -/var/log/secure;RSYSLOG_TraditionalFileFormat
```

The '-' sign in front of the file path is optional.

Remediation:

If the `ActionFileDefaultTemplate` setting is missing in file `/etc/rsyslog.conf`, add it at the top.

```
printf '1,a\n$ActionFileDefaultTemplate RSYSLOG_TraditionalFileFormat\n.\nw\nq\n' | ed $file
```

If the entry for 'messages' is missing in file `/etc/rsyslog.conf`, add it:

```
printf '*.info;mail.none;authpriv.none;cron.none\n/var/log/messages\n' >> $file
```

If the entry for 'secure' is missing in file `/etc/rsyslog.conf`, add it:

```
printf 'authpriv.* /var/log/secure\n' >> $file
```

If you applied any change to `/etc/rsyslog.conf`, trigger a re-read of the configuration file:

```
pkill -HUP rsyslogd
```

Default Value:

N/A

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.			

5 Access, Authentication and Authorization

5.1 Configure cron

5.1.1 Ensure cron daemon is enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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:

Run the the following command to verify `cron` is enabled:

```
# systemctl is-enabled crond  
enabled
```

Verify result is "enabled".

Remediation:

Run the following command to enable `cron`:

```
# systemctl --now enable crond
```

Additional Information:

Additional methods of enabling a service exist. Consult your distribution documentation for appropriate methods.

5.1.2 Ensure permissions on /etc/crontab are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 these files could provide unprivileged users with the ability to elevate their privileges. Read access to these files 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:

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:

Run the following commands to set ownership and permissions on `/etc/crontab` :

```
# chown root:root /etc/crontab
# chmod og-rwx /etc/crontab
```

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.			

5.1.3 Ensure permissions on /etc/cron.hourly are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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:

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:

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
```


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.			

5.1.4 Ensure permissions on /etc/cron.daily are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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:

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:

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
```

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.			

5.1.5 Ensure permissions on /etc/cron.weekly are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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:

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:

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
```

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.			

5.1.6 Ensure permissions on /etc/cron.monthly are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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:

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:

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
```

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>3.3 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	<u>14.6 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.			

5.1.7 Ensure permissions on /etc/cron.d are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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:

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:

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
```

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.			

5.1.8 Ensure at/cron is restricted to authorized users (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Configure `/etc/cron.allow` and `/etc/at.allow` to allow specific users to use these services. If `/etc/cron.allow` or `/etc/at.allow` do not exist, then `/etc/at.deny` and `/etc/cron.deny` are checked. Any user not specifically defined in those files is allowed to use at and cron. By removing the files, only users in `/etc/cron.allow` and `/etc/at.allow` are allowed to use at and cron. Note that 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:

Run the following commands and ensure `/etc/cron.deny` and `/etc/at.deny` do not exist:

```
# stat /etc/cron.deny
stat: cannot stat `/etc/cron.deny': No such file or directory

# stat /etc/at.deny
stat: cannot stat `/etc/at.deny': No such file or directory
```

Run the following command and verify Uid and Gid are both 0/root and Access does not grant permissions to group or other for both `/etc/cron.allow` and `/etc/at.allow`:

```
# stat /etc/cron.allow
Access: (0600/-rw-----)  Uid: (   0/   root)  Gid: (   0/   root)

# stat /etc/at.allow
Access: (0600/-rw-----)  Uid: (   0/   root)  Gid: (   0/   root)
```

Remediation:

Run the following commands to remove `/etc/cron.deny` and `/etc/at.deny` and create and set permissions and ownership for `/etc/cron.allow` and `/etc/at.allow`:

```
# rm /etc/cron.deny
# rm /etc/at.deny

# touch /etc/cron.allow
# touch /etc/at.allow

# chmod og-rwx /etc/cron.allow
# chmod og-rwx /etc/at.allow

# chown root:root /etc/cron.allow
# chown root:root /etc/at.allow
```

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.			

5.2 SSH Server Configuration

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.

Note: Once all configuration changes have been made to `/etc/ssh/sshd_config`, the `sshd` configuration must be reloaded:

```
# systemctl reload sshd
```

5.2.1 Ensure permissions on /etc/ssh/sshd_config are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 /etc/ssh/sshd_config
Access: (0600/-rw-----)  Uid: (   0/   root)   Gid: (   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 og-rwx /etc/ssh/sshd_config
```

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.			

5.2.2 Ensure SSH access is limited (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 command:

```
sshd -T | grep -E '^s*(allow|deny) (users|groups) \s+\S+'
```

Verify that the output matches at least one of the following lines:

```
AllowUsers <userlist>  
AllowGroups <grouplist>  
DenyUsers <userlist>  
DenyGroups <grouplist>
```

Remediation:

Edit the `/etc/ssh/sshd_config` file to set one or more of the parameter as follows:

```
AllowUsers <userlist>  
AllowGroups <grouplist>  
DenyUsers <userlist>  
DenyGroups <grouplist>
```

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.			

5.2.3 Ensure permissions on SSH private host key files are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 command and verify Uid is 0/root and Gid is 0/root. Ensure group and other do not have permissions

```
# find /etc/ssh -xdev -type f -name 'ssh_host_*_key' -exec stat {} \;

  File: '/etc/ssh/ssh_host_rsa_key'
  Size: 1679          Blocks: 8          IO Block: 4096   regular file
Device: ca01h/51713d  Inode: 8628138      Links: 1
Access: (0600/-rw-----)  Uid: (  0/   root)   Gid: (  0/   root)
Access: 2018-10-22 18:24:56.861750616 +0000
Modify: 2018-10-22 18:24:56.861750616 +0000
Change: 2018-10-22 18:24:56.873750616 +0000
Birth: -
  File: '/etc/ssh/ssh_host_ecdsa_key'
  Size: 227          Blocks: 8          IO Block: 4096   regular file
Device: ca01h/51713d  Inode: 8631760      Links: 1
Access: (0600/-rw-----)  Uid: (  0/   root)   Gid: (  0/   root)
Access: 2018-10-22 18:24:56.897750616 +0000
Modify: 2018-10-22 18:24:56.897750616 +0000
Change: 2018-10-22 18:24:56.905750616 +0000
Birth: -
  File: '/etc/ssh/ssh_host_ed25519_key'
  Size: 387          Blocks: 8          IO Block: 4096   regular file
Device: ca01h/51713d  Inode: 8631762      Links: 1
Access: (0600/-rw-----)  Uid: (  0/   root)   Gid: (  0/   root)
Access: 2018-10-22 18:24:56.945750616 +0000
Modify: 2018-10-22 18:24:56.945750616 +0000
Change: 2018-10-22 18:24:56.957750616 +0000
Birth: -
```

Remediation:

Run the following commands to set ownership and permissions on the private SSH host key files

```
# find /etc/ssh -xdev -type f -name 'ssh_host_*_key' -exec chown root:root {} \;
# find /etc/ssh -xdev -type f -name 'ssh_host_*_key' -exec chmod 0600 {} \;
```

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.			

5.2.4 Ensure permissions on SSH public host key files are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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

```
# find /etc/ssh -xdev -type f -name 'ssh_host_*_key.pub' -exec stat {} \;

  File: '/etc/ssh/ssh_host_rsa_key.pub'
  Size: 382          Blocks: 8          IO Block: 4096   regular file
Device: ca01h/51713d Inode: 8631758      Links: 1
Access: (0644/-rw-r--r--)  Uid: (    0/   root)   Gid: (    0/   root)
Access: 2018-10-22 18:24:56.861750616 +0000
Modify: 2018-10-22 18:24:56.861750616 +0000
Change: 2018-10-22 18:24:56.881750616 +0000
Birth: -
  File: '/etc/ssh/ssh_host_ecdsa_key.pub'
  Size: 162          Blocks: 8          IO Block: 4096   regular file
Device: ca01h/51713d Inode: 8631761      Links: 1
Access: (0644/-rw-r--r--)  Uid: (    0/   root)   Gid: (    0/   root)
Access: 2018-10-22 18:24:56.897750616 +0000
Modify: 2018-10-22 18:24:56.897750616 +0000
Change: 2018-10-22 18:24:56.917750616 +0000
Birth: -
  File: '/etc/ssh/ssh_host_ed25519_key.pub'
  Size: 82           Blocks: 8          IO Block: 4096   regular file
Device: ca01h/51713d Inode: 8631763      Links: 1
Access: (0644/-rw-r--r--)  Uid: (    0/   root)   Gid: (    0/   root)
Access: 2018-10-22 18:24:56.945750616 +0000
Modify: 2018-10-22 18:24:56.945750616 +0000
Change: 2018-10-22 18:24:56.961750616 +0000
Birth: -
```

Remediation:

Run the following commands to set permissions and ownership on the SSH host public key files

```
# find /etc/ssh -xdev -type f -name 'ssh_host_*_key.pub' -exec chmod 0644 {} \;
#find /etc/ssh -xdev -type f -name 'ssh_host_*_key.pub' -exec chown root:root {} \;
```

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	14.3 <u>Disable Workstation to Workstation Communication</u> Disable all workstation to workstation communication to limit an attacker's ability to move laterally and compromise neighboring systems, through technologies such as Private VLANs or microsegmentation.			

5.2.5 Ensure SSH LogLevel is appropriate (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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:

```
# sshd -T | grep loglevel
LogLevel VERBOSE

OR

loglevel INFO
```

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the parameter as follows:

```
LogLevel VERBOSE
```

OR

```
LogLevel INFO
```

Default Value:

LogLevel INFO

References:

1. https://www.ssh.com/ssh/sshd_config/

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.			

5.2.6 Ensure SSH X11 forwarding is disabled (Automated)

Profile Applicability:

- Level 1 - Workstation
- 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 | grep x11forwarding
X11Forwarding no
```

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the parameter as follows:

```
X11Forwarding no
```

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.			

5.2.7 Ensure SSH MaxAuthTries is set to 4 or less (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 | grep maxauthtries  
MaxAuthTries 4
```

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the parameter as follows:

```
MaxAuthTries 4
```

Default Value:

MaxAuthTries 6

5.2.8 Ensure SSH IgnoreRhosts is enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 | grep ignorerhosts  
  
IgnoreRhosts yes
```

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the parameter as follows:

```
IgnoreRhosts yes
```

Default Value:

`IgnoreRhosts yes`

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.			

5.2.9 Ensure SSH HostbasedAuthentication is disabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 | grep hostbasedauthentication
HostbasedAuthentication no
```

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the parameter as follows:

```
HostbasedAuthentication no
```

Default Value:

HostbasedAuthentication no

5.2.10 Ensure SSH root login is disabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 | grep permitrootlogin  
PermitRootLogin no
```

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the parameter as follows:

```
PermitRootLogin no
```

Default Value:

`PermitRootLogin without-password`

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.			

5.2.11 Ensure SSH PermitEmptyPasswords is disabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 | grep permitemptypasswords  
PermitEmptyPasswords no
```

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the parameter as follows:

```
PermitEmptyPasswords no
```

Default Value:

`PermitEmptyPasswords no`

5.2.12 Ensure SSH PermitUserEnvironment is disabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 | grep permituserenvironment  
PermitUserEnvironment no
```

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the parameter as follows:

```
PermitUserEnvironment no
```

Default Value:

`PermitUserEnvironment no`

5.2.13 Ensure SSH Idle Timeout Interval is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The two options `ClientAliveInterval` and `ClientAliveCountMax` control the timeout of ssh sessions. When the `ClientAliveInterval` variable is set, ssh sessions that have no activity for the specified length of time are terminated. When the `ClientAliveCountMax` variable is set, `sshd` will send client alive messages at every `ClientAliveInterval` interval. When the number of consecutive client alive messages are sent with no response from the client, the `ssh` session is terminated. For example, if the `ClientAliveInterval` is set to 15 seconds and the `ClientAliveCountMax` is set to 3, the client `ssh` session will be terminated after 45 seconds of idle time.

Rationale:

Having no timeout value associated with a connection could allow an unauthorized user access to another user's `ssh` session (e.g. user walks away from their computer and doesn't lock the screen). Setting a timeout value at least reduces the risk of this happening..

While the recommended setting is 300 seconds (5 minutes), set this timeout value based on site policy. The recommended setting for `ClientAliveCountMax` is 0. In this case, the client session will be terminated after 5 minutes of idle time and no keepalive messages will be sent.

Audit:

Run the following commands and verify `ClientAliveInterval` is between 1 and 300 and `ClientAliveCountMax` is 3 or less:

```
# sshd -T | grep clientaliveinterval
ClientAliveInterval 300

# sshd -T | grep clientalivecountmax
ClientAliveCountMax 0
```


Remediation:

Edit the `/etc/ssh/sshd_config` file to set the parameters according to site policy:

```
ClientAliveInterval 300  
ClientAliveCountMax 0
```

Default Value:

ClientAliveInterval 300

ClientAliveCountMax 0

5.2.14 Ensure SSH LoginGraceTime is set to one minute or less (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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:

```
# sshd -T | grep logingracetime  
LoginGraceTime 60
```

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the parameter as follows:

```
LoginGraceTime 60
```

Default Value:

`LoginGraceTime 120`

5.2.15 Ensure SSH warning banner is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 | grep banner  
Banner /etc/issue.net
```

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the parameter as follows:

```
Banner /etc/issue.net
```

5.2.16 Ensure SSH PAM is enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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(8) as a non-root user.

Audit:

Run the following command and verify that output matches:

```
# sshd -T | grep -i usepam  
usepam yes
```

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the parameter as follows:

```
UsePAM yes
```

Default Value:

usePAM yes

5.2.17 Ensure SSH AllowTcpForwarding is disabled (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

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 | grep -i allowtcpforwarding
AllowTcpForwarding no
```

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the parameter as follows:

```
AllowTcpForwarding no
```

Default Value:

AllowTcpForwarding yes

References:

1. <https://www.ssh.com/ssh/tunneling/example>

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.			

5.2.18 Ensure SSH MaxStartups is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 matches site policy:

```
# sshd -T | grep -i maxstartups
maxstartups 10:30:60
```

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the parameter as follows:

```
maxstartups 10:30:60
```

5.2.19 Ensure SSH MaxSessions is set to 4 or less (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 `sshd` logins and prevent overwhelming the daemon.

Audit:

Run the following command and verify that output `MaxSessions` is 4 or less, or matches site policy:

```
# sshd -T | grep -i maxsessions  
maxsessions 4
```

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the parameter as follows:

```
MaxSessions 4
```


5.2.20 Ensure system-wide crypto policy is not over-ridden (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 '^/s*CRYPTO_POLICY=' /etc/sysconfig/sshd
```

No output should be returned

Remediation:

Run the following commands:

```
# sed -ri "s/^\s*(CRYPTO_POLICY\s*=.*)$/# \1/" /etc/sysconfig/sshd
# systemctl reload sshd
```

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.			

5.3 Configure authselect

Authselect is a utility that simplifies the configuration of user authentication on a CentOS Linux host. 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

Notes:

- **Do not use authselect if your host is part of CentOS Linux Identity Management or Active Directory.** The `ipa-client-install` command, called when joining your host to a CentOS Identity Management domain, takes full care of configuring authentication on your host. Similarly the `realm join` command, called when joining your host to an Active Directory domain, takes full care of configuring authentication on your host.
- 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
```

5.3.1 Create custom authselect profile (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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.

Rationale:

A custom profile is required to customize many of the pam options

Audit:

Run the following command:

```
# authselect current | grep "Profile ID: custom"
Profile ID: custom/<custom profile name>
```

Verify that the custom profile follows local site policy

Remediation:

Run the following command to create a custom authselect profile:

```
# authselect create-profile <custom-profile name> -b <default profile to copy>
```

Example:

```
# authselect create-profile custom-profile -b sssd --symlink-meta
```

5.3.2 Select authselect profile (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

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.

You can create and deploy a custom profile by customizing one of the default profiles, the sssd, winbind, or the nis profile.

Rationale:

When you deploy a profile, the profile is applied to every user logging into the given host

Audit:

Run the following command and verify that the current custom authselect profile follows local site policy:

```
# authselect current
```

Output should be similar to:

```
Profile ID: <custom-profile name>
Enabled features:
- with-sudo
- with-faillock
- without-nullok
```

Remediation:

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. Using authselect on a Red Hat Enterprise Linux host

5.3.3 Ensure authselect includes with-faillock (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 deny consecutive failed authentications. 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 authselect includes the with-faillock feature

```
# authselect current | grep with-faillock
- with-faillock
# grep with-faillock /etc/authselect/authselect.conf
with-faillock
```

Remediation:

Run the following command to include the `with-faillock` option

```
# authselect select <PROFILE NAME> with-faillock
```

Example:

```
# authselect select custom/custom-profile with-sudo with-faillock without-nullok
```

References:

1. faillock(8) - Linux man page

5.4 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.

5.4.1 Ensure password creation requirements are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 to update the system-auth and password-auth files

```
CP=$(authselect current | awk 'NR == 1 {print $3}' | grep custom/)
for FN in system-auth password-auth; do
    [[ -n $CP ]] && PTF=/etc/authselect/$CP/$FN || PTF=/etc/authselect/$FN
    [[ -z $(grep -E
'^\s*password\s+requisite\s+pam_pwquality.so\s+.*enforce-for-root\s*.*$'
$PTF) ]] && sed -ri
's/^\s*(password\s+requisite\s+pam_pwquality.so\s+)(.*)$/\1\2 enforce-for-
root/' $PTF
    [[ -n $(grep -E
'^\s*password\s+requisite\s+pam_pwquality.so\s+.*\s+retry=\S+\s*.*$' $PTF) ]]
&& sed -ri '/pwquality/s/retry=\S+/retry=3/' $PTF || sed -ri
's/^\s*(password\s+requisite\s+pam_pwquality.so\s+)(.*)$/\1\2 retry=3/' $PTF
done
authselect apply-changes
```

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.			

5.4.2 Ensure lockout for failed password attempts is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Lock out users after n unsuccessful consecutive login attempts.

- deny= - Number of attempts before the account is locked
- unlock_time= - Time in seconds before the account is unlocked

Set the lockout number and unlock time to follow local site policy.

Rationale:

Locking out user IDs after n unsuccessful consecutive login attempts mitigates brute force password attacks against your systems.

Audit:

Verify password lockouts are configured. These settings are commonly configured with the `pam_tally2.so` and `pam_faillock.so` modules found in `/etc/pam.d/common-auth` or `/etc/pam.d/system-auth` and `/etc/pam.d/password-auth`. Examples:

Run the following command and review the output to ensure that it follows local site policy. `deny` should be no greater than 5 and `unlock_time` should be no less than 900 seconds

```
# grep -E '^\\s*auth\\s+required\\s+pam_faillock.so\\s+' /etc/pam.d/password-auth
/etc/pam.d/system-auth
```

Output should look similar to:

```
/etc/pam.d/password-auth:auth      required      pam_faillock.so preauth silent
deny=5 unlock_time=900
/etc/pam.d/password-auth:auth      required      pam_faillock.so authfail
deny=5 unlock_time=900
/etc/pam.d/system-auth:auth        required      pam_faillock.so preauth silent
deny=5 unlock_time=900
/etc/pam.d/system-auth:auth        required      pam_faillock.so authfail
deny=5 unlock_time=900
```

Remediation:

Set password lockouts and unlock times to conform to site policy

Run the following to update the `system-auth` and `password-auth` files. This script will update/add the `deny=5` and `unlock_time=900` options.

This script should be modified as needed to follow local site policy.

```
CP=$(authselect current | awk 'NR == 1 {print $3}' | grep custom/)
for FN in system-auth password-auth; do
    [[ -n $CP ]] && PTF=/etc/authselect/$CP/$FN || PTF=/etc/authselect/$FN
    [[ -n $(grep -E
'^\\s*auth\\s+required\\s+pam_faillock.so\\s+.*deny=\\S+\\s*.*$' $PTF) ]] && sed -
ri '/pam_faillock.so/s/deny=\\S+/deny=5/g' $PTF || sed -ri
's/^\\s*(auth\\s+required\\s+pam_faillock\\.so\\s+)(\\.^[^{}]) (\\{\\.\\*\\}|)\\$/\\1\\2
deny=5 \\3/' $PTF
    [[ -n $(grep -E
'^\\s*auth\\s+required\\s+pam_faillock.so\\s+.*unlock_time=\\S+\\s*.*$' $PTF) ]] &&
sed -ri '/pam_faillock.so/s/unlock_time=\\S+/unlock_time=900/g' $PTF || sed -
ri 's/^\\s*(auth\\s+required\\s+pam_faillock\\.so\\s+)(\\.^[^{}]) (\\{\\.\\*\\}|)\\$/\\1\\2
unlock_time=900 \\3/' $PTF
done
authselect apply-changes
```

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 -u --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
v7	16.7 <u>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.			

5.4.3 Ensure password reuse is limited (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 that 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 -E  
'^\s*password\s+(requisite|sufficient)\s+(pam_pwquality\.so|pam_unix\.so)\s+.  
*remember=([5-9]|[1-4][0-9])[0-9]*\s*.*$' /etc/pam.d/system-auth
```

The output should be similar to:

```
password      requisite      pam_pwquality.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:

```
CP=$(authselect current | awk 'NR == 1 {print $3}' | grep custom/)
[[ -n $CP ]] && PTF=/etc/authselect/$CP/system-auth ||
PTF=/etc/authselect/system-auth
[[ -n $(grep -E
'^\s*password\s+(sufficient\s+pam_unix|requi(red|site)\s+pam_pwhistory)\.so\s+
([^\s]+\s+)*remember=\S+\s*.*$' $PTF) ]] && sed -ri
's/^\s*(password\s+(requisite\sufficient)\s+(pam_pwquality\.so|pam_unix\.so)\s+
(.*) (remember=\S+\s*) (.*)$/\1\4 remember=5 \6/' $PTF || sed -ri
's/^\s*(password\s+(requisite\sufficient)\s+(pam_pwquality\.so|pam_unix\.so)\s+
(.*)$/\1\4 remember=5/' $PTF
authselect apply-changes
```

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.			

5.4.4 Ensure password hashing algorithm is SHA-512 (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The commands below change password encryption from `md5` to `sha512` (a much stronger hashing algorithm). All existing accounts will need to perform a password change to upgrade the stored hashes to the new algorithm.

Rationale:

The SHA-512 algorithm provides much stronger hashing than MD5, thus providing additional protection to the system by increasing the level of effort for an attacker to successfully determine passwords.

Note that these changes only apply to accounts configured on the local system.

Audit:

Verify password hashing algorithm is `sha512`. This setting is configured with the `pam_unix.so sha512` option found in `/etc/pam.d/system-auth` and `/etc/pam.d/password-auth`

Run the following command:

```
# grep -E '^s*password\s+sufficient\s+pam_unix.so\s+.*sha512\s*.*$' /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_auth tok
/etc/pam.d/system-auth:password    sufficient    pam_unix.so sha512
shadow  try_first_pass use_auth tok remember=5
```

Remediation:

Set password hashing algorithm to sha512.

Run the following script to add or modify the `pam_unix.so` lines in the `password-auth` and `system-auth` files to include the `sha512` option:

```
CP=$(authselect current | awk 'NR == 1 {print $3}' | grep custom/)
for FN in system-auth password-auth; do
    [[ -z $(grep -E
'^\s*password\s+sufficient\s+pam_unix.so\s+.*sha512\s*.*$' $PTF) ]] && sed -
ri 's/^\s*(password\s+sufficient\s+pam_unix.so\s+)(.*)$/\1\2 sha512/' $PTF
done
authselect apply-changes
```

Additional Information:

Additional module options may be set, recommendation only covers those listed here.

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. To accomplish that, the following commands can be used. 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 Encrypt Sensitive Data at Rest 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 Encrypt or Hash all Authentication Credentials Encrypt or hash with a salt all authentication credentials when stored.			

5.5 User Accounts and Environment

This section provides guidance on setting up secure defaults for system and user accounts and their environment.

5.5.1 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.

5.5.1.1 Ensure password expiration is 365 days or less (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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>
```

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	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.			

5.5.1.2 Ensure minimum days between password changes is 7 or more (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 7 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 7 days):

```
# grep PASS_MIN_DAYS /etc/login.defs  
PASS_MIN_DAYS 7
```

Run the following command and Review list of users and `PASS_MIN_DAYS` to Verify that all users' `PASS_MIN_DAYS` conform s to site policy (no less than 7 days):

```
# grep -E ^[^:]+:[^!*] /etc/shadow | cut -d: -f1,4  
<user>:<PASS_MIN_DAYS>
```

Remediation:

Set the `PASS_MIN_DAYS` parameter to 7 in `/etc/login.defs` :

```
PASS_MIN_DAYS 7
```

Modify user parameters for all users with a password set to match:

```
# chage --mindays 7 <user>
```

Additional Information:

You can also check this setting in `/etc/shadow` directly. The 4th field should be 7 or more for all users with a password.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.2 <u>Use Unique Passwords</u> 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 <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.			

5.5.1.3 Ensure password expiration warning days is 7 or more (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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>
```

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	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.			

5.5.1.4 Ensure inactive password lock is 30 days or less (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

User accounts that have been inactive for over a given period of time can be automatically disabled. It is recommended that accounts that are inactive for 30 days after password expiration be disabled.

Rationale:

Inactive accounts pose a threat to system security since the users are not logging in to notice failed login attempts or other anomalies.

Audit:

Run the following command and verify `INACTIVE` conforms to site policy (no more than 30 days):

```
# useradd -D | grep INACTIVE  
  
INACTIVE=30
```

Verify all users with a password have Password inactive no more than 30 days after password expires:

Run the following command and Review list of users and `INACTIVE` to verify that all users' `INACTIVE` conforms to site policy (no more than 30 days):

```
# grep -E ^[^:]+:[^!*] /etc/shadow | cut -d: -f1,7  
  
<user>:<INACTIVE>
```

Remediation:

Run the following command to set the default password inactivity period to 30 days:

```
# useradd -D -f 30
```

Modify user parameters for all users with a password set to match:

```
# chage --inactive 30 <user>
```

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 <u>Use Unique Passwords</u> 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 <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.			

5.5.1.5 Ensure all users last password change date is in the past (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

All users should have a password change date in the past.

Rationale:

If a users 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

```
# for usr in $(cut -d: -f1 /etc/shadow); do [[ $(chage --list $usr | grep '^Last password change' | cut -d: -f2) > $(date) ]] && echo "$usr :$(chage --list $usr | grep '^Last password change' | cut -d: -f2)"; 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.

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.			

5.5.2 Ensure system accounts are secured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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.

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:

Run the following commands and verify no results are returned:

```
awk -F: '($1!="root" && $1!="sync" && $1!="shutdown" && $1!="halt" && $1!~/^\/ && $3<"$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)"' && $7!="$(which nologin)"' && $7!="/bin/false") {print}' /etc/passwd
awk -F: '($1!="root" && $1!~/^\/ && $3<"$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)"') {print $1}' /etc/passwd | xargs -I '{}' passwd -S '{}' |
awk '($2!="L" && $2!="LK") {print $1}'
```

Remediation:

Run the commands appropriate for your distribution:

Set the shell for any accounts returned by the audit to nologin:

```
# usermod -s $(which nologin) <user>
```

Lock any non root accounts returned by the audit:

```
# usermod -L <user>
```

The following command will set all system accounts to a non login shell:

```
awk -F: '($1!="root" && $1!="sync" && $1!="shutdown" && $1!="halt" && $1~/^\/ && $3<"$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)" && $7!="$(which nologin)" && $7!="bin/false") {print $1}' /etc/passwd | while read user do usermod -s $(which nologin) $user done
```

The following command will automatically lock not root system accounts:

```
awk -F: '($1!="root" && $1~/^\/ && $3<"$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)"') {print $1}' /etc/passwd | xargs -I '{}' passwd -S '{}' | awk '($2!="L" && $2!="LK") {print $1}' | while read user do usermod -L $user done
```

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
v7	16 <u>Account Monitoring and Control</u> Account Monitoring and Control			

5.5.3 Ensure default user shell timeout is 900 seconds or less (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The default `TMOUT` determines the shell timeout for users. The `TMOUT` value is measured in seconds.

Rationale:

Having no timeout value associated with a shell could allow an unauthorized user access to another user's shell session (e.g. user walks away from their computer and doesn't lock the screen). Setting a timeout value at least reduces the risk of this happening.

Audit:

Run the following commands and verify all `TMOUT` lines returned are 900 or less and at least one exists in each file.

```
# grep "^TMOUT" /etc/bashrc
readonly TMOUT=900 ; export TMOUT

# grep "^TMOUT" /etc/profile /etc/profile.d/*.sh
readonly TMOUT=900 ; export TMOUT
```

Remediation:

Edit the `/etc/bashrc`, `/etc/profile` and `/etc/profile.d/*.sh` files (and the appropriate files for any other shell supported on your system) and add or edit any `umask` parameters as follows:

```
readonly TMOUT=900 ; export TMOUT
```

Note that setting the value to `readonly` prevents unwanted modification during runtime.

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	<u>4.3 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	<u>16.11 Lock Workstation Sessions After Inactivity</u> Automatically lock workstation sessions after a standard period of inactivity.			

5.5.4 Ensure default group for the root account is GID 0 (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `usermod` command can be used to specify which group the root user belongs to. This affects permissions of files that are created by the root user.

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` user default group to GID 0 :

```
# usermod -g 0 root
```


5.5.5 Ensure default user umask is 027 or more restrictive (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The default `umask` determines the permissions of files created by users. The user creating the file has the discretion of making their files and directories readable by others via the `chmod` command. Users who wish to allow their files and directories to be readable by others by default may choose a different default `umask` by inserting the `umask` command into the standard shell configuration files (`.profile` , `.bashrc` , etc.) in their home directories.

Rationale:

Setting a very secure default value for `umask` ensures that users make a conscious choice about their file permissions. A default `umask` setting of `077` causes files and directories created by users to not be readable by any other user on the system. A `umask` of `027` would make files and directories readable by users in the same Unix group, while a `umask` of `022` would make files readable by every user on the system.

Audit:

Run the following commands and verify all `umask` lines returned are `027` or more restrictive.

```
# grep "umask" /etc/bashrc
umask 027
# grep "umask" /etc/profile /etc/profile.d/*.sh
umask 027
```

Remediation:

Edit the `/etc/bashrc`, `/etc/profile` and `/etc/profile.d/*.sh` files (and the appropriate files for any other shell supported on your system) and add or edit any `umask` parameters as follows:

```
umask 027
```

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 default user umask exist however the shell configuration files are the last run and will override other settings if they exist therefor our recommendation is to configure in the shell configuration files. If other methods are in use in your environment they should be audited and the shell configs should be verified to not override.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>3.3 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	<u>14.6 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.			

5.6 Ensure root login is restricted to system console (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The file `/etc/securetty` contains a list of valid terminals that may be logged in directly as root.

Rationale:

Since the system console has special properties to handle emergency situations, it is important to ensure that the console is in a physically secure location and that unauthorized consoles have not been defined.

Audit:

```
# cat /etc/securetty
```

Remediation:

Remove entries for any consoles that are not in a physically secure location.

5.7 Ensure access to the su command is restricted (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 the `wheel` group to execute `su`.

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 output includes matching line:

```
# grep pam_wheel.so /etc/pam.d/su
auth required pam_wheel.so use_uid
```

Run the following command and verify users in `wheel` group match site policy. If no users are listed, only root will have access to `su`.

```
# grep wheel /etc/group
wheel:x:10:root,<user list>
```

Remediation:

Add the following line to the `/etc/pam.d/su` file:

```
auth required pam_wheel.so use_uid
```

Create a comma separated list of users in the wheel statement in the `/etc/group` file:

```
wheel:x:<GID>:root,<user list>
```

Example:

```
wheel:x:10:root,user1,user2,user3
```

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.			

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 Audit system file permissions (Manual)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

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
# dpkg -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
# dpkg --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

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 storable 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 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.			

6.1.2 Ensure permissions on /etc/passwd are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 and verify `Uid` and `Gid` are both `0/root` and Access is `644` :

```
# stat /etc/passwd
Access: (0644/-rw-r--r--)  Uid: (    0/   root)   Gid: (    0/   root)
```

Remediation:

Run the following command to set permissions on `/etc/passwd` :

```
# chown root:root /etc/passwd
# chmod 644 /etc/passwd
```

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.			

6.1.3 Ensure permissions on /etc/shadow are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 and verify `Uid` is `0/root`, `Gid` is `0/root` or `<gid>/shadow`, and `Access` is `640` or more restrictive:

```
# stat /etc/shadow
Access: (0640/-rw-r-----)  Uid: (   0/   root)  Gid: (   0/   root)
```

Remediation:

Run one of the following `chown` commands as appropriate and the `chmod` to set permissions on `/etc/shadow`:

```
# chown root:root /etc/shadow
# chown root:shadow /etc/shadow

# chmod o-rwx,g-wx /etc/shadow
```

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.			

6.1.4 Ensure permissions on /etc/group are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 and verify `Uid` and `Gid` are both `0/root` and Access is `644` :

```
# stat /etc/group
Access: (0644/-rw-r--r--)  Uid: (   0/   root)   Gid: (   0/   root)
```

Remediation:

Run the following command to set permissions on `/etc/group` :

```
# chown root:root /etc/group
# chmod 644 /etc/group
```

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.			

6.1.5 Ensure permissions on /etc/gshadow are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 and verify `Uid` is `0/root`, `Gid` is `0/root` or `<gid>/shadow`, and `Access` is `640` or more restrictive:

```
# stat /etc/gshadow
Access: (0640/-rw-r-----)  Uid: (   0/   root)   Gid: (   0/   root)
```

Remediation:

Run one of the following `chown` commands as appropriate and the `chmod` to set permissions on `/etc/gshadow`:

```
# chown root:root /etc/gshadow
# chown root:shadow /etc/gshadow

# chmod o-rwx,g-rw /etc/gshadow
```

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.			

6.1.6 Ensure permissions on /etc/passwd- are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 and verify Uid and Gid are both 0/root and Access is 600 or more restrictive:

```
# stat /etc/passwd-  
Access: (0600/-rw-----)  Uid: (    0/    root)   Gid: (    0/    root)
```

Remediation:

Run the following command to set permissions on /etc/passwd- :

```
# chown root:root /etc/passwd-  
# chmod u-x,go-rwx /etc/passwd-
```

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.			

6.1.7 Ensure permissions on /etc/shadow- are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 and verify `Uid` is `0/root`, `Gid` is `0/root` or `<gid>/shadow`, and `Access` is `640` or more restrictive:

```
# stat /etc/shadow-  
Access: (0600/-rw-----)  Uid: (    0/    root)  Gid: (    0/    root)
```

Remediation:

Run the one of the following `chown` commands as appropriate and the `chmod` to set permissions on `/etc/shadow-` :

```
# chown root:root /etc/shadow-
```

OR

```
# chown root:shadow /etc/shadow-  
# chmod u-x,go-rwx /etc/shadow-
```

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.			

6.1.8 Ensure permissions on /etc/group- are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 and verify Uid and Gid are both 0/root and Access is 644 or more restrictive:

```
# stat /etc/group-  
Access: (0644/-rw-----)  Uid: (    0/    root)   Gid: (    0/    root)
```

Remediation:

Run the following command to set permissions on /etc/group- :

```
# chown root:root /etc/group-  
# chmod u-x,go-wx /etc/group-
```

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.			

6.1.9 Ensure permissions on /etc/gshadow- are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 and verify `Uid` is `0/root`, `Gid` is `0/root` or `<gid>/shadow`, and `Access` is `640` or more restrictive:

```
# stat /etc/gshadow-  
Access: (0640/-rw-r-----)  Uid: (    0/    root)  Gid: (    0/    root)
```

Remediation:

Run one of the following `chown` commands as appropriate and the `chmod` to set permissions on `/etc/gshadow-` :

```
# chown root:root /etc/gshadow-  
# chown root:shadow /etc/gshadow-  
  
# chmod o-rwx,g-rw /etc/gshadow-
```

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>3.3 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	<u>14.6 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.			

6.1.10 Ensure no world writable files exist (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Unix-based systems support variable settings to control access to files. World writable files are the least secure. See the `chmod(2)` man page for more information.

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.

Audit:

Run the following command and verify no files are returned:

```
# df --local -P | awk '{if (NR!=1) print $6}' | xargs -I '{}' find '{}' -xdev -type f -perm -0002
```

The command above only searches local filesystems, there may still be compromised items on network mounted partitions. Additionally the `--local` option to `df` is not universal to all versions, it can be omitted to search all filesystems on a system including network mounted filesystems or the following command can be run manually for each partition:

```
# find <partition> -xdev -type f -perm -0002
```

Remediation:

Removing write access for the "other" category (`chmod o-w <filename>`) is advisable, but always consult relevant vendor documentation to avoid breaking any application dependencies on a given 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.			

6.1.11 Ensure no unowned files or directories exist (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Sometimes when administrators delete users from the password file they neglect to remove all files owned by those users from the system.

Rationale:

A new user who is assigned the deleted user's user ID or group ID may then end up "owning" these files, and thus have more access on the system than was intended.

Audit:

Run the following command and verify no files are returned:

```
# df --local -P | awk {'if (NR!=1) print $6'} | xargs -I '{}' find '{}' -xdev -nouser
```

The command above only searches local filesystems, there may still be compromised items on network mounted partitions. Additionally the `--local` option to `df` is not universal to all versions, it can be omitted to search all filesystems on a system including network mounted filesystems or the following command can be run manually for each partition:

```
# find <partition> -xdev -nouser
```

Remediation:

Locate files that are owned by users or groups not listed in the system configuration files, and reset the ownership of these files to some active user on the system as appropriate.

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.			

6.1.12 Ensure no ungrouped files or directories exist (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Sometimes when administrators delete users or groups from the system they neglect to remove all files owned by those users or groups.

Rationale:

A new user who is assigned the deleted user's user ID or group ID may then end up "owning" these files, and thus have more access on the system than was intended.

Audit:

Run the following command and verify no files are returned:

```
# df --local -P | awk '{if (NR!=1) print $6}' | xargs -I '{}' find '{}' -xdev -nogroup
```

The command above only searches local filesystems, there may still be compromised items on network mounted partitions. Additionally the `--local` option to `df` is not universal to all versions, it can be omitted to search all filesystems on a system including network mounted filesystems or the following command can be run manually for each partition:

```
# find <partition> -xdev -nogroup
```

Remediation:

Locate files that are owned by users or groups not listed in the system configuration files, and reset the ownership of these files to some active user on the system as appropriate.

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.			

6.1.13 Audit SUID executables (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 program is to enable users to perform functions (such as changing their password) that require root privileges.

Rationale:

There are valid reasons for SUID programs, but it is important to identify and review such programs to ensure they are legitimate.

Audit:

Run the following command to list SUID files:

```
# df --local -P | awk '{if (NR!=1) print $6}' | xargs -I '{}' find '{}' -xdev -type f -perm -4000
```

The command above only searches local filesystems, there may still be compromised items on network mounted partitions. Additionally the `--local` option to `df` is not universal to all versions, it can be omitted to search all filesystems on a system including network mounted filesystems or the following command can be run manually for each partition:

```
# find <partition> -xdev -type f -perm -4000
```

Remediation:

Ensure that no rogue SUID 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.

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.			

6.1.14 Audit SGID executables (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 SGID program is to enable users to perform functions (such as changing their password) that require root privileges.

Rationale:

There are valid reasons for 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 md5 checksum than what from the package. This is an indication that the binary may have been replaced.

Audit:

Run the following command to list SGID files:

```
# df --local -P | awk '{if (NR!=1) print $6}' | xargs -I '{}' find '{}' -xdev -type f -perm -2000
```

The command above only searches local filesystems, there may still be compromised items on network mounted partitions. Additionally the `--local` option to `df` is not universal to all versions, it can be omitted to search all filesystems on a system including network mounted filesystems or the following command can be run manually for each partition:

```
# find <partition> -xdev -type f -perm -2000
```

Remediation:

Ensure that no rogue 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.

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.			

6.1.15 Ensure udev rules have proper ownership (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Udev rules are triggered by kernel events and are executed in **root** context. It is critical that existing rules can only be updated by a **root** user.

Rationale:

If a regular user would be able or allowed to update udev rules, that specific user could execute arbitrary code in **root** context. This is most probably not desired.

Audit:

Check the following files and directories:

```
ls -l /usr/lib/udev
ls -l /etc/udev/rules.d
```

Make sure that only the **root** user has write-access to the directory and existing files.

Remediation:

If the udev rules do not have the proper **root** user ownership, change the access rights by using the following commands:

```
chown -R root.root /usr/lib/udev
chown -R root.root /etc/udev/rules.d
chmod -R go-w /usr/lib/udev
chmod -R go-w /etc/udev/rules.d
```

Default Value:

root.root

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.			

6.2 User and Group Settings

This section provides guidance on securing aspects of the 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 password fields are not empty (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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.

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.			

6.2.2 Ensure no legacy "+" entries exist in /etc/passwd (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The character + in various files used to be markers for systems to insert data from NIS maps at a certain point in a system configuration file. These entries are no longer required on most systems, but may exist in files that have been imported from other platforms.

Rationale:

These entries may provide an avenue for attackers to gain privileged access on the system.

Audit:

Run the following command and verify that no output is returned:

```
# grep '^\\+: ' /etc/passwd
```

Remediation:

Remove any legacy '+' entries from /etc/passwd if they exist.

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.			

6.2.3 Ensure root PATH Integrity (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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:

```
for x in $(echo $PATH | tr ":" " "); do
    if [ -d "$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.

6.2.4 Ensure no legacy "+" entries exist in /etc/shadow (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The character + in various files used to be markers for systems to insert data from NIS maps at a certain point in a system configuration file. These entries are no longer required on most systems, but may exist in files that have been imported from other platforms.

Rationale:

These entries may provide an avenue for attackers to gain privileged access on the system.

Audit:

Run the following command and verify that no output is returned:

```
# grep '^\\+: ' /etc/shadow
```

Remediation:

Remove any legacy '+' entries from `/etc/shadow` if they exist.

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.			

6.2.5 Ensure no legacy "+" entries exist in /etc/group (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The character + in various files used to be markers for systems to insert data from NIS maps at a certain point in a system configuration file. These entries are no longer required on most systems, but may exist in files that have been imported from other platforms.

Rationale:

These entries may provide an avenue for attackers to gain privileged access on the system.

Audit:

Run the following command and verify that no output is returned:

```
# grep '^\\+: ' /etc/group
```

Remediation:

Remove any legacy '+' entries from /etc/group if they exist.

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.			

6.2.6 Ensure root is the only UID 0 account (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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.

6.2.7 Ensure users' home directories permissions are 750 or more restrictive (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

While the system administrator can establish secure permissions for users' home directories, the users can easily override these.

Rationale:

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.

Audit:

Run the following script and verify no results are returned:

```
#!/bin/bash
grep -E -v '^(halt|sync|shutdown)' /etc/passwd | awk -F: '($7 != ""$(which
nologin)"" && $7 != "/bin/false") { print $1 " " $6 }' | while read user
dir; do
    if [ ! -d "$dir" ]; then
        echo "The home directory ($dir) of user $user does not exist."
    else
        dirperm=$(ls -ld $dir | cut -f1 -d" ")
        if [ $(echo $dirperm | cut -c6) != "-" ]; then
            echo "Group Write permission set on the home directory ($dir) of user
$user"
        fi
        if [ $(echo $dirperm | cut -c8) != "-" ]; then
            echo "Other Read permission set on the home directory ($dir) of user
$user"
        fi
        if [ $(echo $dirperm | cut -c9) != "-" ]; then
            echo "Other Write permission set on the home directory ($dir) of user
$user"
        fi
        if [ $(echo $dirperm | cut -c10) != "-" ]; then
            echo "Other Execute permission set on the home directory ($dir) of user
$user"
        fi
    fi
done
```

Remediation:

Making global modifications to user home directories 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 file permissions and determine the action to be taken in accordance with site policy.

Additional Information:

On some distributions the `/sbin/nologin` should be replaced with `/usr/sbin/nologin`.

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.			

6.2.8 Ensure users own their home directories (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The user home directory is space defined for the particular user to set local environment variables and to store personal files.

Rationale:

Since the user is accountable for files stored in the user home directory, the user must be the owner of the directory.

Audit:

Run the following script and verify no results are returned:

```
#!/bin/bash

grep -E -v '^(halt|sync|shutdown)' /etc/passwd | awk -F: '($7 != ""$(which nologin)"" && $7 != "/bin/false") { print $1 " " $6 }' | while read user
dir; do
    if [ ! -d "$dir" ]; then
        echo "The home directory ($dir) of user $user does not exist."
    else
        owner=$(stat -L -c "%U" "$dir")
        if [ "$owner" != "$user" ]; then
            echo "The home directory ($dir) of user $user is owned by $owner."
        fi
    fi
done
```

Remediation:

Change the ownership of any home directories that are not owned by the defined user to the correct user.

Additional Information:

On some distributions the `/sbin/nologin` should be replaced with `/usr/sbin/nologin`.

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.			

6.2.9 Ensure users' dot files are not group or world writable (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

While the system administrator can establish secure permissions for users' "dot" files, the users can easily override these.

Rationale:

Group or world-writable user configuration files may enable malicious users to steal or modify other users' data or to gain another user's system privileges.

Audit:

Run the following script and verify no results are returned:

```
#!/bin/bash

grep -E -v '^(halt|sync|shutdown)' /etc/passwd | awk -F: '($7 != ""$(which nologin)"" && $7 != "/bin/false") { print $1 " " $6 }' | while read user dir; do
    if [ ! -d "$dir" ]; then
        echo "The home directory ($dir) of user $user does not exist."
    else
        for file in $dir/.[A-Za-z0-9]*; do
            if [ ! -h "$file" -a -f "$file" ]; then
                fileperm=$(ls -ld $file | cut -f1 -d" ")

                if [ $(echo $fileperm | cut -c6) != "-" ]; then
                    echo "Group Write permission set on file $file"
                fi
                if [ $(echo $fileperm | cut -c9) != "-" ]; then
                    echo "Other Write permission set on file $file"
                fi
            fi
        done
    fi
done
```

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.

Additional Information:

On some distributions the `/sbin/nologin` should be replaced with `/usr/sbin/nologin`.

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.			

6.2.10 Ensure no users have .forward files (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `.forward` file specifies an email address to forward the user's mail to.

Rationale:

Use of the `.forward` file poses a security risk in that sensitive data may be inadvertently transferred outside the organization. The `.forward` file also poses a risk as it can be used to execute commands that may perform unintended actions.

Audit:

Run the following script and verify no results are returned:

```
#!/bin/bash

grep -E -v '^(root|halt|sync|shutdown)' /etc/passwd | awk -F: '($7 !=
'"$(which nologin)"' && $7 != "/bin/false") { print $1 " " $6 }' | while
read user dir; do
    if [ ! -d "$dir" ]; then
        echo "The home directory ($dir) of user $user does not exist."
    else
        if [ ! -h "$dir/.forward" -a -f "$dir/.forward" ]; then
            echo ".forward file $dir/.forward exists"
        fi
    fi
done
```

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 `.forward` files and determine the action to be taken in accordance with site policy.

Additional Information:

On some distributions the `/sbin/nologin` should be replaced with `/usr/sbin/nologin`.

6.2.11 Ensure no users have .netrc files (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `.netrc` file contains data for logging into a remote host for file transfers via FTP.

Rationale:

The `.netrc` file presents a significant security risk since it stores passwords in unencrypted form. Even if FTP is disabled, user accounts may have brought over `.netrc` files from other systems which could pose a risk to those systems.

Audit:

Run the following script and verify no results are returned:

```
#!/bin/bash

grep -E -v '^(root|halt|sync|shutdown)' /etc/passwd | awk -F: '($7 !=
'"$(which nologin)"' && $7 != "/bin/false") { print $1 " " $6 }' | while
read user dir; do
    if [ ! -d "$dir" ]; then
        echo "The home directory ($dir) of user $user does not exist."
    else
        if [ ! -h "$dir/.netrc" -a -f "$dir/.netrc" ]; then
            echo ".netrc file $dir/.netrc exists"
        fi
    fi
done
```

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 `.netrc` files and determine the action to be taken in accordance with site policy.

Additional Information:

On some distributions the `/sbin/nologin` should be replaced with `/usr/sbin/nologin`.

6.2.12 Ensure users' .netrc Files are not group or world accessible (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

While the system administrator can establish secure permissions for users' `.netrc` files, the users can easily override these.

Rationale:

`.netrc` files may contain unencrypted passwords that may be used to attack other systems.

Audit:

Run the following script and verify no results are returned:

```
#!/bin/bash

grep -E -v '^(root|halt|sync|shutdown)' /etc/passwd | awk -F: '($7 !=
""$(which nologin)"" && $7 != "/bin/false") { print $1 " " $6 }' | while
read user dir; do
    if [ ! -d "$dir" ]; then
        echo "The home directory ($dir) of user $user does not exist."
    else
        for file in $dir/.netrc; do
            if [ ! -h "$file" -a -f "$file" ]; then
                fileperm=$(ls -ld $file | cut -f1 -d" ")
                if [ $(echo $fileperm | cut -c5) != "-" ]; then
                    echo "Group Read set on $file"
                fi
                if [ $(echo $fileperm | cut -c6) != "-" ]; then
                    echo "Group Write set on $file"
                fi
                if [ $(echo $fileperm | cut -c7) != "-" ]; then
                    echo "Group Execute set on $file"
                fi
                if [ $(echo $fileperm | cut -c8) != "-" ]; then
                    echo "Other Read set on $file"
                fi
                if [ $(echo $fileperm | cut -c9) != "-" ]; then
                    echo "Other Write set on $file"
                fi
                if [ $(echo $fileperm | cut -c10) != "-" ]; then
                    echo "Other Execute set on $file"
                fi
            fi
        done
    fi
done
```

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 `.netrc` file permissions and determine the action to be taken in accordance with site policy.

Additional Information:

While the complete removal of `.netrc` files is recommended if any are required on the system secure permissions must be applied.

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.			

6.2.13 Ensure no users have .rhosts files (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

While no `.rhosts` files are shipped by default, users can easily create them.

Rationale:

This action is only meaningful if `.rhosts` support is permitted in the file `/etc/pam.conf`. Even though the `.rhosts` files are ineffective if support is disabled in `/etc/pam.conf`, they may have been brought over from other systems and could contain information useful to an attacker for those other systems.

Audit:

Run the following script and verify no results are returned:

```
#!/bin/bash

grep -E -v '^(root|halt|sync|shutdown)' /etc/passwd | awk -F: '($7 !=
'"$(which nologin)"' && $7 != "/bin/false") { print $1 " " $6 }' | while
read user dir; do
    if [ ! -d "$dir" ]; then
        echo "The home directory ($dir) of user $user does not exist."
    else
        for file in $dir/.rhosts; do
            if [ ! -h "$file" -a -f "$file" ]; then
                echo ".rhosts file in $dir"
            fi
        done
    fi
done
```

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 `.rhosts` files and determine the action to be taken in accordance with site policy.

Additional Information:

On some distributions the `/sbin/nologin` should be replaced with `/usr/sbin/nologin`.

6.2.14 Ensure all groups in /etc/passwd exist in /etc/group (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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.

6.2.15 Ensure no duplicate UIDs exist (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.2 <u>Use Unique Passwords</u> 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 <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.			

6.2.16 Ensure no duplicate GIDs exist (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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.

Additional Information:

You can also use the `grpck` command to check for other inconsistencies in the `/etc/group` file.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.2 <u>Use Unique Passwords</u> 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 <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.			

6.2.17 Ensure no duplicate user names exist (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 user name.

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 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.

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.			

6.2.18 Ensure no duplicate group names exist (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

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 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.

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.			

6.2.19 Ensure shadow group is empty (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The shadow group allows system programs which require access the ability to read the `/etc/shadow` file. No users should be assigned to the shadow group.

Rationale:

Any users assigned to the shadow group would be granted read access to the `/etc/shadow` file. If attackers can gain read access to the `/etc/shadow` file, they can easily run a password cracking program against the hashed passwords to break them. Other security information that is stored in the `/etc/shadow` file (such as expiration) could also be useful to subvert additional user accounts.

Audit:

Run the following commands and verify no results are returned:

```
# grep ^shadow:[^:]*:[^:]*:[^:]+ /etc/group
# awk -F: '($4 == "<shadow-gid>") { print }' /etc/passwd
```

Remediation:

Remove all users from the shadow group, and change the primary group of any users with shadow as their primary group.

6.2.20 Ensure all users' home directories exist (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Users can be defined in `/etc/passwd` without a home directory or with a home directory that does not actually exist.

Rationale:

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 and verify no results are returned:

```
#!/bin/bash
grep -E -v '^(halt|sync|shutdown)' /etc/passwd | awk -F: '($7 != "" || $(which
nologin) "" && $7 != "/bin/false") { print $1 " " $6 }' | while read -r user
dir; do
    if [ ! -d "$dir" ]; then
        echo "The home directory ($dir) of user $user does not exist."
    fi
done
```

Remediation:

If any users' home directories do not exist, create them and make sure the respective user owns the directory. Users without an assigned home directory should be removed or assigned a home directory as appropriate.

Additional Information:

The audit script checks all users with interactive shells except `halt`, `sync`, `shutdown`, and `nfsnobody`.

6.3 Enable timeout support for login shells (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

When establishing an individual to a system ID, a login shell which supports time out must be used. The guidance from other sections in this benchmark is only sufficient for login shells: `/bin/csh`, `/bin/tcsh`, `/bin/sh`, `/bin/ksh`, `/bin/ksh93`, and `/bin/bash`. This guidance is also sufficient for `/bin/sh` when the concerned link is a symbolic or hard link to one of the previously listed shells.

If other login shells are permitted, similar controls must be established. Where this is done, add it to the 'agreed to Value' column. Neither the bourne shell nor the dash shell support this control, and neither shell may be used as a login shell.

Shells from `/bin/false`, `/sbin/nologin`, and other no-login shells are exempt from this requirement and are permitted. Additional no-login shells defined by the account can be added to the 'agreed to Value' column. The special-purpose shells from `/bin/sync`, `/sbin/shutdown`, and `/sbin/halt` are also exempt from this requirement, but must not be set as the login shell for an individual to a system ID.

An individual logging in to system IDs must have one of the following shells set as login shell: `/bin/csh`, `/bin/tcsh`, `/bin/sh`, `/bin/ksh`, `/bin/bash`, `/bin/sh`, `/bin/false`, `/sbin/nologin`, `/usr/bin/sh`, `/usr/bin/bash`, `/usr/sbin/nologin`, `/bin/ksh93`, or `/usr/bin/ksh93`.

Rationale:

Having no timeout value associated with a shell could allow an unauthorized user access to another user's shell session (for example, the user walks away from their computer and does not lock the screen). Setting a timeout value at least reduces this risk.

Audit:

Run the following procedure:

```
file=/etc/passwd
Shells='/bin/csh|/bin/tcsh|/bin/sh|/bin/ksh|/bin/bash|/bin/sh|/bin/false|
        /sbin/nologin|/usr/bin/sh|/usr/bin/bash|/usr/sbin/nologin|/bin
        /ksh93|/usr/bin/ksh|/usr/bin/rksh|/usr/bin/ksh93|'
cat $file | egrep -v '^sync:./bin/sync\s$|^halt:./sbin/halt\s$|^shutdown:./
        :/sbin/shutdown\s$' |
while read ENTRY
do
SHELL=$(echo $ENTRY | awk -F: '{printf $7"\n"}')
Uid=$(echo $ENTRY | awk -F: '{printf $1"\n"}')
[[ "$SHELL" != @($Shells) ]] && printf "
WARNING: Verify shell supports autologout/TMOUT function - id:
$Uid - shell: $SHELL\n"
done
```

This procedure should return nothing.

Note: This procedure also passes without having defined shells, as Linux assigns an effective default shell `/bin/sh` to the ID.

Remediation:

Check if the indicated shell supports time out.

If it does and the customer requires it's usage, update the technical specification 'agreed to values' to reflect this requirement. Be sure to update the 'default' requirements for setting timeout for the shell. That is, set systems to terminate a session after a period of inactivity as follows:

PCI Servers = 15 minutes, or 900 seconds.

GDPR Servers = 30 minutes, or 1800 seconds.

HIPAA Servers = 30 minutes, or 1800 seconds.

All other servers = 6 hours, or 21600 seconds.

If the shell is to be permitted to the general population, add it to the shells in file

`/etc/shells.`

```
usermod --shell /bin/bash $Uid
```

Default Value:

N/A

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	<u>4.3 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	<u>16.11 Lock Workstation Sessions After Inactivity</u> Automatically lock workstation sessions after a standard period of inactivity.			

7 Hardware

Recommendations in this section cover multiple aspects related to hardware.

7.1 Cryptography

This subsection provides guidance on security aspects concerning your system hardware.

7.1.1 Ensure correct assignment of cryptographic coprocessors and domains for EP11 tokens (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Ensure that for each openCryptoki EP11 token the correct set of cryptographic coprocessors and domains to be used by the token is assigned.

Rationale:

An incorrect configuration can lead to the usage of a wrong cryptographic coprocessor for cryptographic operations.

Audit:

Check the set of assigned cryptographic coprocessor and domains for each available openCryptoki EP11 token in the appropriate EP11 token-specific configuration files.

Remediation:

Specify the correct assignment of cryptographic coprocessors and domains for each EP11 token as described in [Defining an EP11 token-specific configuration file](#).

Default Value:

N/A

7.1.2 Ensure correct control domain assignment (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

When managing cryptographic devices on the Trusted Key Entry (TKE) workstation, the operating system needs access to control domains. Ensure that control domains are only assigned to LPARs that are allowed to manage master keys within a hardware security module (HSM).

Rationale:

Unauthorized change of HSM master keys can weaken or disable cryptographic operations performed by other systems.

Audit:

On the Hardware Management Console (HMC), make sure that control domains are only assigned to LPARs that need to manage cryptographic devices via the TKE workstation.

Remediation:

Change the assignment of control domains via the HMC (Hardware Management Console). See [Hardware Management Console Operations Guide](#). Download the PDF and navigate to topic: *Tasks -> HMC Tasks -> Change LPAR Cryptographic Controls*.

Default Value:

n/a

7.1.3 Ensure access to cryptographic devices is allowed only through device nodes created by zcryptctl (Manual)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

In containerized environments, ensure that access to cryptographic devices is only allowed through multi device node extensions created by

```
zcryptctl
```

with the correct assignment of cryptographic coprocessors and domains.

Rationale:

Assigning all host cryptographic devices to a container allows the containerized workload unwanted arbitrary access to cryptographic resources.

Audit:

Enter the following command:

```
zcryptctl list
```

Ensure that each container is assigned a specific device node with the correct assignment of cryptographic coprocessors and domains.

Remediation:

For each container create additional crypto devices as described in [zcryptctl - Control access to AP queues and functions](#)

Default Value:

N/A. Additional cryptographic devices must be created explicitly.

7.1.4 Configure redundant APQNs and make sure they are on different adapters (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

When performing cryptographic operations, it is recommended to provide redundant cryptographic resources (that is, APQNs) on different physical cryptographic adapters.

Rationale:

If the cryptographic coprocessors (adapters) are not redundant, the system is more susceptible to hardware failure. Also, the ability to perform hardware-supported cryptographic operations at all may get lost.

Audit:

Enter the following command:

```
lszcrypt
CARD.DOMAIN TYPE  MODE          STATUS  REQUESTS
-----
08           CEX7C CCA-Coproc   online    4
08.0031      CEX7C CCA-Coproc   online    4
09           CEX7C CCA-Coproc   online    2
09.0031      CEX7C CCA-Coproc   online    2
0a           CEX7P EP11-Coproc  online    0
0a.0031      CEX7P EP11-Coproc  online    0
```

Verify that APQNs are available on different adapters (different values in the **CARD** column).

Remediation:

Configure additional APQNs for the system by adding adapters and domains the the LPAR on the HMC (Hardware Management Console). See [Hardware Management Console Operations Guide](#). Download the PDF and navigate to topic: *Tasks -> HMC Tasks -> Change LPAR Cryptographic Controls*.

Default Value:

n/a

7.1.5 Enable CPACF feature (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

It is recommended to enable the CP Assist for Cryptographic Functions (CPACF) feature when performing supported cryptographic operations.

Rationale:

Without CPACF enabled, cryptographic operations for supported algorithms are not accelerated by the hardware, thus reducing overall workload performance.

Audit:

Enter the following command:

```
# icainfo
```

In the resulting output, make sure that for the value in column **hardware > static** reads *yes* for the relevant cryptographic algorithms, as shown in the following sample excerpt:

Cryptographic algorithm support				
function	hardware			software
	dynamic	static		
SHA-1	no	yes		yes
SHA-224	no	yes		yes
...				
SHA-512/224	no	yes		yes
SHA-512/256	no	yes		yes
...				
AES ECB	no	yes		yes
AES CBC	no	yes		yes
AES OFB	no	yes		no
...				

Remediation:

Enable CPACF on the Hardware Management Console (HMC). See [Hardware Management Console Operations Guide](#). Download the PDF and navigate to topic: *Tasks -> HMC Tasks -> Customize/Delete Activation Profiles -> CP Assist for Cryptographic Functions*.

Default Value:

n/a

7.2 I/O

This subsection provides guidance on input/output aspects concerning your system hardware.

7.2.1 Ensure cio_ignore is configured (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

When a Linux on Z instance boots, it senses and analyzes all available I/O devices. You can use the `cio_ignore` command to specify devices that are to be ignored. So ensure that `cio_ignore` is used to only make required channel devices available.

Rationale:

Using devices inadvertently can lead to data corruption or leaks.

Audit:

Enter the following command to display specifications for all devices that are not in the current exclusion list, that is, they are currently not ignored:

```
cio_ignore -L
```

Make sure that only devices are listed that are required for the current workload. Additionally enter the following command:
`grep cio_ignore=all /proc/cmdline`
In the resulting output, verify that all devices are ignored by default on system boot.
For example:

```
# grep cio_ignore=all /proc/cmdline
root=/dev/disk/by-path/ccw-0.0.1c80-part1
crashkernel=128M rd.dasd=0.0.1c80 rd.dasd=0.0.1c83
cio_ignore=all``,!condev rd.znet=qeth,0.0.bdf0,0.0.bdf1,
0.0.bdf2,layer2=1,portno=0 BOOT_IMAGE=0
```

Remediation:

Add `cio_ignore=all` to the kernel command line in the boot loader entry in `/boot/loader/entries`.
Then run the `zipl` command and restart the system.

Default Value:

Devices not required for system boot are ignored by default.

7.2.2 Enable SMART monitor (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Ensure that the SMART monitor is enabled to be notified of disk issues in order to prevent silent disk failures.

Rationale:

Without monitoring SMART events for disk devices, errors leading to disk failures can not be identified and disk devices can fail before being replaced.

Audit:

Enter the following command to ensure that the smartd.service is enabled and running:

```
systemctl status smartd.service
```

Remediation:

To enable the SMART monitor service, enter the following command:

```
systemctl enable --now smartd.service
```

Default Value:

enabled

7.2.3 Configure dm-crypt to use XTS-PAES with 256 bit keys (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

When using full disk encryption via `dm-crypt`, it is recommended to use the XTS-PAES cipher mode with 256 bit keys.

Rationale:

A key length of 256 bits is required to ensure future-proof quantum-safe encryption. Using the XTS-PAES cipher mode results in a much smaller passphrase making key management easier for the user.

Audit:

Enter the following command to verify the cipher mode:

```
$ cryptsetup luksDump <device> | grep -E "^Cipher mode|^MK bits"
Cipher mode:      xts-paes
MK bits:          256
```

Remediation:

Reformat the device using
`cryptsetup luksFormat`
and specify
`--cipher xts-paes`
and `--key-size 256`.

See also:

[Pervasive Encryption](#)

Default Value:

See `cryptsetup --help`.

7.3 Peripheral Component Interconnect (PCI) Devices

This subsection provides guidance on security aspects concerning Peripheral Component Interconnect (PCI) devices .

7.3.1 Ensure UID uniqueness checking is active (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Activate UID uniqueness checking to ensure stable addressing for Peripheral Component Interconnect (PCI) devices.

Rationale:

Deactivating or activating a PCI device can change its corresponding PCI domain, resulting in a change of the NIC (network interface cardname) device in the operating system. This can compromise security for mechanisms that rely on the NIC, like for example, `iptables`.

Audit:

Read the ***uid_is_unique*** attribute for any PCIe device that is available to your Linux instance to find out which naming scheme applies.

Example:

```
# cat /sys/bus/pci/devices/0000:00:00.0/uid_is_unique
1
```

If the value is 1, UID uniqueness checking is enabled, and the network interface names are based on UIDs. For any other value, UID uniqueness checking is not enabled, and the network interface names are based on FIDs.

Remediation:

UID uniqueness checking can be enabled in the IOCDs.

Default Value:

- 0:** for classic mode LPAR
- 1:** for DPM partitions

7.3.2 NVMe setting for data resilience (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

When using NVMe (Non-Volatile Memory express) devices for disk storage, it is recommended to implement redundancy over multiple devices via software RAID for data resilience.

Rationale:

NVMe devices do not have built-in hardware capabilities for redundancy similar to classic IO devices (DASD, zfcP). Therefore, redundancy needs to be implemented at the software level. Without redundancy, disk failure of a single disk can lead to data loss.

Audit:

To check software RAID at block device level, enter the following command:

`lsblk`

In the output, verify that the mount point of the specific file system is distributed across redundant physical devices, for example:

NAME	MAJ:MIN	RM	SIZE	RO	TYPE	MOUNTPOINT
dasda	94:0	0	20.6G	0	disk	
`-dasda1	94:1	0	20.6G	0	part	/
dasdb	94:4	0	20.6G	0	disk	
`-dasdb1	94:5	0	20.6G	0	part	[SWAP]

To check software RAID at file system level, use the utilities specific for the used file system to show information about the RAID status, for example, for the `Btrfs` (B-Tree filesystem):

```
$ btrfs filesystem show /mnt/  
Label: none  uuid: 46c3b60c-9abe-41b9-9d4e-efd3da0c059a  
Total devices 2 FS bytes used <size>  
devid    1 size <path> used <size> path <path>  
devid    2 size <path> used <size> path <path>
```

In the resulting output, verify that the file system is distributed across redundant physical devices.

Remediation:

Configure software RAID across redundant devices either via block device level or file system level mechanisms.

Default Value:

n/a

7.3.3 Ensure that unused PCI Devices are disabled (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

In general it is recommended to keep unused PCI (Peripheral Component Interconnect) devices disabled so that for example, NICs (Network Interface Card) are not used inadvertently.

Rationale:

Inadvertent usage of devices can lead to data corruption or leakage.

Audit:

Enter the following command:

```
lspci
```

Make sure that only devices are listed in the resulting output, that are required for the current workload.

Remediation:

Enter the following command to detach the devices in the IOCDs:

```
echo 0 >/sys/bus/pci/slots/<slot>/power
```

This ensures that any events are no longer processed by the effected hardware.

Default Value:

n/a

7.4 Network

Recommendations in this section cover multiple aspects related to the network to which your system is connected.

7.4.1 Configure SMC-D for network traffic isolation (Manual)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

When using Shared Memory Communication - Direct Memory Access (SMC-D) networking, it is recommended that each SMC-D device belongs to a different virtual channel ID (VCHID).

Rationale:

All SMC-D devices belonging to the same VCHID might be able to eavesdrop traffic.

Audit:

Use the **smc_rnics** tool to list, enable and disable (R)NICs as used by SMC-R and SMC-D.

Make sure that each SMC-D device has a unique physical channel identifier (PCHID) value.

Remediation:

Configure individual SMC-D devices to have different VCHIDs.

Default Value:

n/a

7.4.2 Ensure RoCE card redundancy settings are configured properly (Manual)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

When using channel-bonding over multiple RoCE NICs, it is recommended that each PCI function ID (FID) belongs to a different support partition on the CEC.

Rationale:

If the support partition is shared between PCI FIDs in the same bond, and that support partition becomes unavailable (for example, by an upgrade procedure), the whole channel bond becomes unavailable.

Audit:

Enter the following command:

```
cat /sys/bus/pci/devices/<addr>/pfip/segment0
```

Make sure that the value for the PCI support partition is different for all RoCE devices.

Remediation:

Support partitions for PCI FIDs can only be configured via changes in the IOCDs data set.

Default Value:

n/a

Appendix: Summary Table

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1	Initial Setup		
1.1	Filesystem Configuration		
1.1.1	Disable unused filesystems		
1.1.1.1	Ensure mounting of cramfs filesystems is disabled (Automated)	☐	☐
1.1.1.2	Ensure mounting of vFAT filesystems is limited (Manual)	☐	☐
1.1.1.3	Ensure mounting of squashfs filesystems is disabled (Automated)	☐	☐
1.1.1.4	Ensure mounting of udf filesystems is disabled (Automated)	☐	☐
1.1.2	Ensure /tmp is configured (Automated)	☐	☐
1.1.3	Ensure nodev option set on /tmp partition (Automated)	☐	☐
1.1.4	Ensure nosuid option set on /tmp partition (Automated)	☐	☐
1.1.5	Ensure noexec option set on /tmp partition (Automated)	☐	☐
1.1.6	Ensure separate partition exists for /var (Automated)	☐	☐
1.1.7	Ensure separate partition exists for /var/tmp (Automated)	☐	☐
1.1.8	Ensure nodev option set on /var/tmp partition (Automated)	☐	☐
1.1.9	Ensure nosuid option set on /var/tmp partition (Automated)	☐	☐
1.1.10	Ensure noexec option set on /var/tmp partition (Automated)	☐	☐
1.1.11	Ensure separate partition exists for /var/log (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.1.12	Ensure separate partition exists for /var/log/audit (Automated)	☐	☐
1.1.13	Ensure separate partition exists for /home (Automated)	☐	☐
1.1.14	Ensure nodev option set on /home partition (Automated)	☐	☐
1.1.15	Ensure nodev option set on /dev/shm partition (Automated)	☐	☐
1.1.16	Ensure nosuid option set on /dev/shm partition (Automated)	☐	☐
1.1.17	Ensure noexec option set on /dev/shm partition (Automated)	☐	☐
1.1.18	Ensure nodev option set on removable media partitions (Manual)	☐	☐
1.1.19	Ensure nosuid option set on removable media partitions (Manual)	☐	☐
1.1.20	Ensure noexec option set on removable media partitions (Manual)	☐	☐
1.1.21	Ensure sticky bit is set on all world-writable directories (Automated)	☐	☐
1.1.22	Disable Automounting (Automated)	☐	☐
1.1.23	Disable USB Storage (Automated)	☐	☐
1.1.24	Mount devices read-only that do not require write access (Optional) (Manual)	☐	☐
1.2	Configure Software Updates		
1.2.1	Ensure GPG keys are configured (Manual)	☐	☐
1.2.2	Ensure gpgcheck is globally activated (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.2.3	Ensure package manager repositories are configured (Manual)	☐	☐
1.3	Configure sudo		
1.3.1	Ensure sudo is installed (Automated)	☐	☐
1.3.2	Ensure sudo commands use pty (Automated)	☐	☐
1.3.3	Ensure sudo log file exists (Automated)	☐	☐
1.3.4	Ensure sudo-specific activity logging (Manual)	☐	☐
1.4	Filesystem Integrity Checking		
1.4.1	Ensure AIDE is installed (Automated)	☐	☐
1.4.2	Ensure filesystem integrity is regularly checked (Automated)	☐	☐
1.5	Secure Boot Settings		
1.5.1	Ensure a secure IPL configuration (Manual)	☐	☐
1.5.2	Ensure permissions on the boot loader configuration are configured (Automated)	☐	☐
1.5.3	Ensure boot loader password is set (Automated)	☐	☐
1.5.4	Ensure authentication required for single user mode (Automated)	☐	☐
1.5.5	Enable secure boot as IPL parameter (Manual)	☐	☐
1.6	Additional Process Hardening		
1.6.1	Ensure core dumps are restricted (Automated)	☐	☐
1.6.2	Ensure address space layout randomization (ASLR) is enabled (Automated)	☐	☐
1.7	Mandatory Access Control		

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.7.1	Configure SELinux		
1.7.1.1	Ensure SELinux is installed (Automated)	☐	☐
1.7.1.2	Ensure SELinux is not disabled in boot loader configuration (Automated)	☐	☐
1.7.1.3	Ensure SELinux policy is configured (Automated)	☐	☐
1.7.1.4	Ensure the SELinux state is set to 'enforcing' (Automated)	☐	☐
1.7.1.5	Ensure no unconfined services exist (Automated)	☐	☐
1.7.1.6	Ensure SETroubleshoot is not installed (Automated)	☐	☐
1.7.1.7	Ensure the MCS Translation Service (mcstrans) is not installed (Automated)	☐	☐
1.7.1.8	Ensure SELinux is not disabled in zipl boot loader (Manual)	☐	☐
1.8	Warning Banners		
1.8.1	Command Line Warning Banners		
1.8.1.1	Ensure message of the day is configured properly (Automated)	☐	☐
1.8.1.2	Ensure local login warning banner is configured properly (Automated)	☐	☐
1.8.1.3	Ensure remote login warning banner is configured properly (Automated)	☐	☐
1.8.1.4	Ensure permissions on /etc/motd are configured (Automated)	☐	☐
1.8.1.5	Ensure permissions on /etc/issue are configured (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.8.1.6	Ensure permissions on /etc/issue.net are configured (Automated)	☐	☐
1.8.2	Ensure GDM login banner is configured (Automated)	☐	☐
1.9	Ensure updates, patches, and additional security software are installed (Manual)	☐	☐
1.10	Ensure system-wide crypto policy is not legacy (Automated)	☐	☐
1.11	Ensure system-wide crypto policy is FUTURE or FIPS (Automated)	☐	☐
1.12	Ensure that the kernel has loaded all s390-specific crypto modules (Manual)	☐	☐
2	Services		
2.1	inetd Services		
2.1.1	Ensure xinetd is not installed (Automated)	☐	☐
2.2	Special Purpose Services		
2.2.1	Time Synchronization		
2.2.1.1	Ensure time synchronization is in use (Manual)	☐	☐
2.2.1.2	Ensure chrony is configured (Automated)	☐	☐
2.2.2	Ensure X Window System is not installed (Automated)	☐	☐
2.2.3	Ensure rsync service is not enabled (Automated)	☐	☐
2.2.4	Ensure Avahi Server is not enabled (Automated)	☐	☐
2.2.5	Ensure SNMP Server is not enabled (Automated)	☐	☐
2.2.6	Disable the pcnfsd service (Automated)	☐	☐
2.2.7	Disable the tcpspray service (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
2.2.8	Ensure HTTP Proxy Server is not enabled (Automated)	☐	☐
2.2.9	Ensure Samba is not enabled (Automated)	☐	☐
2.2.10	Ensure IMAP and POP3 server is not enabled (Automated)	☐	☐
2.2.11	Ensure HTTP server is not enabled (Automated)	☐	☐
2.2.12	Ensure FTP Server is not enabled (Automated)	☐	☐
2.2.13	Ensure DNS Server is not enabled (Automated)	☐	☐
2.2.14	Ensure NFS is not enabled (Automated)	☐	☐
2.2.15	Ensure RPC is not enabled (Automated)	☐	☐
2.2.16	Ensure LDAP server is not enabled (Automated)	☐	☐
2.2.17	Ensure DHCP Server is not enabled (Automated)	☐	☐
2.2.18	Ensure CUPS is not enabled (Automated)	☐	☐
2.2.19	Ensure NIS Server is not enabled (Automated)	☐	☐
2.2.20	Ensure mail transfer agent is configured for local-only mode (Automated)	☐	☐
2.3	Service Clients		
2.3.1	Ensure NIS Client is not installed (Automated)	☐	☐
2.3.2	Ensure telnet client is not installed (Automated)	☐	☐
2.3.3	Ensure LDAP client is not installed (Automated)	☐	☐
3	Network Configuration		
3.1	Network Parameters (Host Only)		
3.1.1	Ensure IP forwarding is disabled (Automated)	☐	☐
3.1.2	Ensure packet redirect sending is disabled (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
3.1.3	Enable SMC-D pre-load libraries (Manual)	☐	☐
3.1.4	Enable VEPA mode traffic for OSA network devices (Manual)	☐	☐
3.1.5	Enable traffic isolation between OSA cards (Manual)	☐	☐
3.2	Network Parameters (Host and Router)		
3.2.1	Ensure source routed packets are not accepted (Automated)	☐	☐
3.2.2	Ensure ICMP redirects are not accepted (Automated)	☐	☐
3.2.3	Ensure secure ICMP redirects are not accepted (Automated)	☐	☐
3.2.4	Ensure suspicious packets are logged (Automated)	☐	☐
3.2.5	Ensure broadcast ICMP requests are ignored (Automated)	☐	☐
3.2.6	Ensure bogus ICMP responses are ignored (Automated)	☐	☐
3.2.7	Ensure reverse path filtering is enabled (Automated)	☐	☐
3.2.8	Ensure TCP SYN Cookies is enabled (Automated)	☐	☐
3.2.9	Ensure IPv6 router advertisements are not accepted (Automated)	☐	☐
3.3	Uncommon Network Protocols		
3.3.1	Ensure DCCP is disabled (Automated)	☐	☐
3.3.2	Disable forwarding of packets (Automated)	☐	☐
3.3.3	Ensure SCTP is disabled (Automated)	☐	☐
3.3.4	Ensure RDS is disabled (Automated)	☐	☐
3.3.5	Ensure TIPC is disabled (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
3.4	Firewall Configuration		
3.4.1	Ensure Firewall software is installed		
3.4.1.1	Ensure a Firewall package is installed (Automated)	☐	☐
3.4.2	Configure firewalld		
3.4.2.1	Ensure firewalld service is enabled and running (Automated)	☐	☐
3.4.2.2	Ensure iptables is not enabled (Automated)	☐	☐
3.4.2.3	Ensure nftables is not enabled (Automated)	☐	☐
3.4.2.4	Ensure default zone is set (Automated)	☐	☐
3.4.2.5	Ensure network interfaces are assigned to appropriate zone (Manual)	☐	☐
3.4.2.6	Ensure unnecessary services and ports are not accepted (Manual)	☐	☐
3.4.3	Configure nftables		
3.4.3.1	Ensure iptables are flushed (Manual)	☐	☐
3.4.3.2	Ensure a table exists (Automated)	☐	☐
3.4.3.3	Ensure base chains exist (Automated)	☐	☐
3.4.3.4	Ensure loopback traffic is configured (Automated)	☐	☐
3.4.3.5	Ensure outbound and established connections are configured (Manual)	☐	☐
3.4.3.6	Ensure default deny firewall policy (Automated)	☐	☐
3.4.3.7	Ensure nftables service is enabled (Automated)	☐	☐
3.4.3.8	Ensure nftables rules are permanent (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
3.4.4	Configure iptables		
3.4.4.1	Configure software		
3.4.4.1.1	Ensure iptables packages are installed (Automated)	☐	☐
3.4.4.1.2	Ensure nftables is not installed (Automated)	☐	☐
3.4.4.1.3	Ensure firewalld is not installed or stopped and masked (Automated)	☐	☐
3.4.4.2	Configure IPv4 iptables		
3.4.4.2.1	Ensure default deny firewall policy (Automated)	☐	☐
3.4.4.2.2	Ensure loopback traffic is configured (Automated)	☐	☐
3.4.4.2.3	Ensure outbound and established connections are configured (Manual)	☐	☐
3.4.4.2.4	Ensure firewall rules exist for all open ports (Automated)	☐	☐
3.4.4.3	Configure IPv6 ip6tables		
3.4.4.3.1	Ensure IPv6 default deny firewall policy (Automated)	☐	☐
3.4.4.3.2	Ensure IPv6 loopback traffic is configured (Automated)	☐	☐
3.4.4.3.3	Ensure IPv6 outbound and established connections are configured (Manual)	☐	☐
3.4.4.3.4	Ensure IPv6 firewall rules exist for all open ports (Manual)	☐	☐
3.5	Ensure wireless interfaces are disabled (Automated)	☐	☐
3.6	Disable IPv6 (Manual)	☐	☐
3.7	Choose hardware-accelerated cipher suites for applications using SSL/TLS (Manual)	☐	☐
4	Logging and Auditing		

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
4.1	Configure System Accounting (auditd)		
4.1.1	Ensure auditing is enabled		
4.1.1.1	Ensure auditd is installed (Automated)	☐	☐
4.1.1.2	Ensure auditd service is enabled (Automated)	☐	☐
4.1.1.3	Enable auditing for processes that start prior to auditd (Automated)	☐	☐
4.1.1.4	Extend the audit backlog limit for the audit daemon (Automated)	☐	☐
4.1.2	Configure Data Retention		
4.1.2.1	Ensure audit log storage size is configured (Automated)	☐	☐
4.1.2.2	Ensure audit logs are not automatically deleted (Automated)	☐	☐
4.1.2.3	Ensure system is disabled when audit logs are full (Automated)	☐	☐
4.1.3	Ensure changes to system administration scope (sudoers) is collected (Automated)	☐	☐
4.1.4	Ensure login and logout events are collected (Automated)	☐	☐
4.1.5	Ensure session initiation information is collected (Automated)	☐	☐
4.1.6	Ensure events that modify date and time information are collected (Automated)	☐	☐
4.1.7	Ensure events that modify the system's Mandatory Access Controls are collected (Automated)	☐	☐
4.1.8	Ensure events that modify the system's network environment are collected (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
4.1.9	Ensure discretionary access control permission modification events are collected (Automated)	☐	☐
4.1.10	Ensure unsuccessful unauthorized file access attempts are collected (Automated)	☐	☐
4.1.11	Ensure events that modify user/group information are collected (Automated)	☐	☐
4.1.12	Ensure successful file system mounts are collected (Automated)	☐	☐
4.1.13	Ensure use of privileged commands is collected (Automated)	☐	☐
4.1.14	Ensure file deletion events by users are collected (Automated)	☐	☐
4.1.15	Ensure kernel module loading and unloading is collected (Automated)	☐	☐
4.1.16	Ensure system administrator actions (sudolog) are collected (Automated)	☐	☐
4.1.17	Ensure the audit configuration is immutable (Automated)	☐	☐
4.2	Configure Logging		
4.2.1	Configure rsyslog		
4.2.1.1	Ensure rsyslog is installed (Automated)	☐	☐
4.2.1.2	Ensure rsyslog Service is enabled (Automated)	☐	☐
4.2.1.3	Ensure rsyslog default file permissions configured (Automated)	☐	☐
4.2.1.4	Ensure logging is configured (Manual)	☐	☐
4.2.1.5	Ensure rsyslog is configured to send logs to a remote log host (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
4.2.1.6	Ensure remote rsyslog messages are only accepted on designated log hosts. (Manual)	☐	☐
4.2.2	Configure journald		
4.2.2.1	Ensure journald is configured to send logs to rsyslog (Automated)	☐	☐
4.2.2.2	Ensure journald is configured to compress large log files (Automated)	☐	☐
4.2.2.3	Ensure journald is configured to write logfiles to persistent disk (Automated)	☐	☐
4.2.3	Ensure permissions on all logfiles are configured (Automated)	☐	☐
4.2.4	Ensure permissions on operating system resources (OSR) (Automated)	☐	☐
4.3	Ensure logrotate is configured (Manual)	☐	☐
4.4	Ensure logging of login success or failure (Automated)	☐	☐
5	Access, Authentication and Authorization		
5.1	Configure cron		
5.1.1	Ensure cron daemon is enabled (Automated)	☐	☐
5.1.2	Ensure permissions on /etc/crontab are configured (Automated)	☐	☐
5.1.3	Ensure permissions on /etc/cron.hourly are configured (Automated)	☐	☐
5.1.4	Ensure permissions on /etc/cron.daily are configured (Automated)	☐	☐
5.1.5	Ensure permissions on /etc/cron.weekly are configured (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
5.1.6	Ensure permissions on /etc/cron.monthly are configured (Automated)	☐	☐
5.1.7	Ensure permissions on /etc/cron.d are configured (Automated)	☐	☐
5.1.8	Ensure at/cron is restricted to authorized users (Automated)	☐	☐
5.2	SSH Server Configuration		
5.2.1	Ensure permissions on /etc/ssh/sshd_config are configured (Automated)	☐	☐
5.2.2	Ensure SSH access is limited (Automated)	☐	☐
5.2.3	Ensure permissions on SSH private host key files are configured (Automated)	☐	☐
5.2.4	Ensure permissions on SSH public host key files are configured (Automated)	☐	☐
5.2.5	Ensure SSH LogLevel is appropriate (Automated)	☐	☐
5.2.6	Ensure SSH X11 forwarding is disabled (Automated)	☐	☐
5.2.7	Ensure SSH MaxAuthTries is set to 4 or less (Automated)	☐	☐
5.2.8	Ensure SSH IgnoreRhosts is enabled (Automated)	☐	☐
5.2.9	Ensure SSH HostbasedAuthentication is disabled (Automated)	☐	☐
5.2.10	Ensure SSH root login is disabled (Automated)	☐	☐
5.2.11	Ensure SSH PermitEmptyPasswords is disabled (Automated)	☐	☐
5.2.12	Ensure SSH PermitUserEnvironment is disabled (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
5.2.13	Ensure SSH Idle Timeout Interval is configured (Automated)	☐	☐
5.2.14	Ensure SSH LoginGraceTime is set to one minute or less (Automated)	☐	☐
5.2.15	Ensure SSH warning banner is configured (Automated)	☐	☐
5.2.16	Ensure SSH PAM is enabled (Automated)	☐	☐
5.2.17	Ensure SSH AllowTcpForwarding is disabled (Automated)	☐	☐
5.2.18	Ensure SSH MaxStartups is configured (Automated)	☐	☐
5.2.19	Ensure SSH MaxSessions is set to 4 or less (Automated)	☐	☐
5.2.20	Ensure system-wide crypto policy is not over-ridden (Automated)	☐	☐
5.3	Configure authselect		
5.3.1	Create custom authselect profile (Automated)	☐	☐
5.3.2	Select authselect profile (Automated)	☐	☐
5.3.3	Ensure authselect includes with-faillock (Automated)	☐	☐
5.4	Configure PAM		
5.4.1	Ensure password creation requirements are configured (Automated)	☐	☐
5.4.2	Ensure lockout for failed password attempts is configured (Automated)	☐	☐
5.4.3	Ensure password reuse is limited (Automated)	☐	☐
5.4.4	Ensure password hashing algorithm is SHA-512 (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
5.5	User Accounts and Environment		
5.5.1	Set Shadow Password Suite Parameters		
5.5.1.1	Ensure password expiration is 365 days or less (Automated)	☐	☐
5.5.1.2	Ensure minimum days between password changes is 7 or more (Automated)	☐	☐
5.5.1.3	Ensure password expiration warning days is 7 or more (Automated)	☐	☐
5.5.1.4	Ensure inactive password lock is 30 days or less (Automated)	☐	☐
5.5.1.5	Ensure all users last password change date is in the past (Automated)	☐	☐
5.5.2	Ensure system accounts are secured (Automated)	☐	☐
5.5.3	Ensure default user shell timeout is 900 seconds or less (Automated)	☐	☐
5.5.4	Ensure default group for the root account is GID 0 (Automated)	☐	☐
5.5.5	Ensure default user umask is 027 or more restrictive (Automated)	☐	☐
5.6	Ensure root login is restricted to system console (Manual)	☐	☐
5.7	Ensure access to the su command is restricted (Automated)	☐	☐
6	System Maintenance		
6.1	System File Permissions		
6.1.1	Audit system file permissions (Manual)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
6.1.2	Ensure permissions on /etc/passwd are configured (Automated)	☐	☐
6.1.3	Ensure permissions on /etc/shadow are configured (Automated)	☐	☐
6.1.4	Ensure permissions on /etc/group are configured (Automated)	☐	☐
6.1.5	Ensure permissions on /etc/gshadow are configured (Automated)	☐	☐
6.1.6	Ensure permissions on /etc/passwd- are configured (Automated)	☐	☐
6.1.7	Ensure permissions on /etc/shadow- are configured (Automated)	☐	☐
6.1.8	Ensure permissions on /etc/group- are configured (Automated)	☐	☐
6.1.9	Ensure permissions on /etc/gshadow- are configured (Automated)	☐	☐
6.1.10	Ensure no world writable files exist (Automated)	☐	☐
6.1.11	Ensure no unowned files or directories exist (Automated)	☐	☐
6.1.12	Ensure no ungrouped files or directories exist (Automated)	☐	☐
6.1.13	Audit SUID executables (Manual)	☐	☐
6.1.14	Audit SGID executables (Manual)	☐	☐
6.1.15	Ensure udev rules have proper ownership (Manual)	☐	☐
6.2	User and Group Settings		
6.2.1	Ensure password fields are not empty (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
6.2.2	Ensure no legacy "+" entries exist in /etc/passwd (Automated)	☐	☐
6.2.3	Ensure root PATH Integrity (Automated)	☐	☐
6.2.4	Ensure no legacy "+" entries exist in /etc/shadow (Automated)	☐	☐
6.2.5	Ensure no legacy "+" entries exist in /etc/group (Automated)	☐	☐
6.2.6	Ensure root is the only UID 0 account (Automated)	☐	☐
6.2.7	Ensure users' home directories permissions are 750 or more restrictive (Automated)	☐	☐
6.2.8	Ensure users own their home directories (Automated)	☐	☐
6.2.9	Ensure users' dot files are not group or world writable (Automated)	☐	☐
6.2.10	Ensure no users have .forward files (Automated)	☐	☐
6.2.11	Ensure no users have .netrc files (Automated)	☐	☐
6.2.12	Ensure users' .netrc Files are not group or world accessible (Automated)	☐	☐
6.2.13	Ensure no users have .rhosts files (Automated)	☐	☐
6.2.14	Ensure all groups in /etc/passwd exist in /etc/group (Automated)	☐	☐
6.2.15	Ensure no duplicate UIDs exist (Automated)	☐	☐
6.2.16	Ensure no duplicate GIDs exist (Automated)	☐	☐
6.2.17	Ensure no duplicate user names exist (Automated)	☐	☐
6.2.18	Ensure no duplicate group names exist (Automated)	☐	☐
6.2.19	Ensure shadow group is empty (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
6.2.20	Ensure all users' home directories exist (Automated)	☐	☐
6.3	Enable timeout support for login shells (Automated)	☐	☐
7	Hardware		
7.1	Cryptography		
7.1.1	Ensure correct assignment of cryptographic coprocessors and domains for EP11 tokens (Manual)	☐	☐
7.1.2	Ensure correct control domain assignment (Manual)	☐	☐
7.1.3	Ensure access to cryptographic devices is allowed only through device nodes created by zcryptctl (Manual)	☐	☐
7.1.4	Configure redundant APQNs and make sure they are on different adapters (Manual)	☐	☐
7.1.5	Enable CPACF feature (Manual)	☐	☐
7.2	I/O		
7.2.1	Ensure cio_ignore is configured (Manual)	☐	☐
7.2.2	Enable SMART monitor (Manual)	☐	☐
7.2.3	Configure dm-crypt to use XTS-PAES with 256 bit keys (Manual)	☐	☐
7.3	Peripheral Component Interconnect (PCI) Devices		
7.3.1	Ensure UID uniqueness checking is active (Manual)	☐	☐
7.3.2	NVMe setting for data resilience (Manual)	☐	☐
7.3.3	Ensure that unused PCI Devices are disabled (Manual)	☐	☐
7.4	Network		
7.4.1	Configure SMC-D for network traffic isolation (Manual)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
7.4.2	Ensure RoCE card redundancy settings are configured properly (Manual)	☐	☐

Appendix: Change History

Date	Version	Changes for this version
May 27, 2022	1.0.0	Published