CIS Center for Internet Security®

CIS Benchmarks™

CIS Ubuntu Linux 26.04 LTS Benchmark

v1.0.0 - 08-26-2026

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/

For information on referencing and/or citing CIS Benchmarks in 3rd party documentation (including using portions of Benchmark Recommendations) please contact CIS Legal (legalnotices@cisecurity.org) and request guidance on copyright usage.

NOTE: It is ***NEVER*** acceptable to host a CIS Benchmark in ***ANY*** format (PDF, etc.) on a 3rd party (non-CIS owned) site.

Internal Only - General

Table of Contents

Terms of Use *1*
Table of Contents *2*
Overview *12*
Important Usage Information **12**
Key Stakeholders **12**
Apply the Correct Version of a Benchmark **13**
Exceptions **13**
Remediation **14**
Summary **14**
Target Technology Details **15**
Intended Audience **15**
Consensus Guidance **16**
Typographical Conventions **17**
Recommendation Definitions *18*
Title **18**
Assessment Status **18**
Automated **18**
Manual **18**
Profile **18**
Description **18**
Rationale Statement **18**
Impact Statement **19**
Audit Procedure **19**
Remediation Procedure **19**
Default Value **19**
References **19**
CIS Critical Security Controls® (CIS Controls®) **19**
Additional Information **19**
Profile Definitions **20**
Acknowledgements **21**
Recommendations *22*
1 Initial Setup **22**
1.1 Filesystem **22**
1.1.1 Configure Filesystem Kernel Modules **23**
1.1.1.1 Ensure cramfs kernel module is not available (Automated) 24
1.1.1.2 Ensure freevxfs kernel module is not available (Automated) 27

Internal Only - General

1.1.1.3 Ensure hfs kernel module is not available (Automated) 30
1.1.1.4 Ensure hfsplus kernel module is not available (Automated) 33
1.1.1.5 Ensure jffs2 kernel module is not available (Automated) 36
1.1.1.6 Ensure overlay kernel module is not available (Automated) 39
1.1.1.7 Ensure squashfs kernel module is not available (Automated) 42
1.1.1.8 Ensure udf kernel module is not available (Automated) 45
1.1.1.9 Ensure firewire-core kernel module is not available (Automated) 48
1.1.1.10 Ensure usb-storage kernel module is not available (Automated) 51
lsmod | grep -P -- 'usb(_|-)storage' 52
modprobe --showconfig | grep -P -- '\b(install|blacklist)\h+usb(_|-)storage\b' 52
1.1.1.11 Ensure unused filesystems kernel modules are not available (Manual) 54
1.1.2 Configure Filesystem Partitions 60
1.1.2.1 Configure /tmp 62
1.1.2.1.1 Ensure /tmp is tmpfs or a separate partition (Automated) 63
1.1.2.1.2 Ensure nodev option set on /tmp partition (Automated) 67
1.1.2.1.3 Ensure nosuid option set on /tmp partition (Automated) 69
1.1.2.1.4 Ensure noexec option set on /tmp partition (Automated) 71
1.1.2.2 Configure /dev/shm 73
1.1.2.2.1 Ensure /dev/shm is tmpfs or a separate partition (Automated) 74
1.1.2.2.2 Ensure nodev option set on /dev/shm partition (Automated) 76
1.1.2.2.3 Ensure nosuid option set on /dev/shm partition (Automated) 78
1.1.2.2.4 Ensure noexec option set on /dev/shm partition (Automated) 80
1.1.2.3 Configure /home 82
1.1.2.3.1 Ensure separate partition exists for /home (Automated) 83
1.1.2.3.2 Ensure nodev option set on /home partition (Automated) 85
1.1.2.3.3 Ensure nosuid option set on /home partition (Automated) 87
1.1.2.4 Configure /var 89
1.1.2.4.1 Ensure separate partition exists for /var (Automated) 90
1.1.2.4.2 Ensure nodev option set on /var partition (Automated) 92
1.1.2.4.3 Ensure nosuid option set on /var partition (Automated) 94
1.1.2.5 Configure /var/tmp 96
1.1.2.5.1 Ensure separate partition exists for /var/tmp (Automated) 97
1.1.2.5.2 Ensure nodev option set on /var/tmp partition (Automated) 99
1.1.2.5.3 Ensure nosuid option set on /var/tmp partition (Automated) 101
1.1.2.5.4 Ensure noexec option set on /var/tmp partition (Automated) 103
1.1.2.6 Configure /var/log 105
1.1.2.6.1 Ensure separate partition exists for /var/log (Automated) 106
1.1.2.6.2 Ensure nodev option set on /var/log partition (Automated) 108
1.1.2.6.3 Ensure nosuid option set on /var/log partition (Automated) 110
1.1.2.6.4 Ensure noexec option set on /var/log partition (Automated) 112
1.1.2.7 Configure /var/log/audit 114
1.1.2.7.1 Ensure separate partition exists for /var/log/audit (Automated) 115
1.1.2.7.2 Ensure nodev option set on /var/log/audit partition (Automated) 117
1.1.2.7.3 Ensure nosuid option set on /var/log/audit partition (Automated) 119
1.1.2.7.4 Ensure noexec option set on /var/log/audit partition (Automated) 121
1.2 Package Management 122
1.2.1 Configure Package Repositories 123
1.2.1.1 Ensure the source.list and .source files use the Signed-By option (Manual) 124
1.2.1.2 Ensure weak dependencies are configured (Automated) 127
1.2.1.3 Ensure access to gpg key files are configured (Automated) 129
1.2.1.4 Ensure access to /etc/apt/trusted.gpg.d directory is configured (Automated) 132
1.2.1.5 Ensure access to /etc/apt/auth.conf.d directory is configured (Automated) 134
1.2.1.6 Ensure access to files in the /etc/apt/auth.conf.d/ directory is configured (Automated) 136
1.2.1.7 Ensure access to /usr/share/keyrings directory is configured (Automated) 138
1.2.1.8 Ensure access to /etc/apt/sources.list.d directory is configured (Automated) 140

Internal Only - General

1.2.1.9 Ensure access to files in /etc/apt/sources.list.d are configured (Automated) 142
1.2.1.10 Ensure repositories are set to https in /etc/apt/sources.list (Automated) 144
1.2.1.11 Ensure repositories are set to https in /etc/apt/sources.list.d directory (Automated) 146
1.2.1.12 Ensure AllowInsecureRepositories is configured (Automated) 148
1.2.1.13 Ensure AllowWeakRepositories is configured (Automated) 150
1.2.1.14 Ensure AllowDowngradeToInsecureRepositories is configured (Automated) 152
1.2.1.15 Ensure Acquire::Check-Date is configured (Automated) 154
1.2.2 Configure Package Updates 156
1.2.2.1 Ensure updates, patches, and additional security software are installed (Manual) ... 157
1.3 Mandatory Access Control 159
1.3.1 Configure AppArmor 160
1.3.1.1 Ensure apparmor packages are installed (Automated) 161
1.3.1.2 Ensure AppArmor is enabled (Automated) 163
1.3.1.3 Ensure all AppArmor Profiles are enforcing (Automated) 165
1.3.1.4 Ensure kernel.apparmor_restrict_unprivileged_unconfined is enabled (Automated). 169
1.4 Configure Bootloader 172
1.4.1 Ensure bootloader password is set (Automated) 173
1.4.2 Ensure access to bootloader config is configured (Automated) 176
1.5 Configure Additional Process Hardening 178
1.5.1 Ensure fs.protected_hardlinks is configured (Automated) 180
1.5.2 Ensure fs.protected_symlinks is configured (Automated) 184
1.5.3 Ensure kernel.yama.ptrace_scope is configured (Automated) 188
1.5.4 Ensure fs.suid_dumpable is configured (Automated) 193
1.5.5 Ensure kernel.dmesg_restrict is configured (Automated) 197
1.5.6 Ensure prelink is not installed (Automated) 201
1.5.7 Ensure Automatic Error Reporting is configured (Automated) 203
1.5.8 Ensure kernel.kptr_restrict is configured (Automated) 205
1.5.9 Ensure kernel.randomize_va_space is configured (Automated) 209
1.5.10 Ensure core file size is configured (Automated) 213
1.5.11 Ensure systemd-coredump ProcessSizeMax is configured (Automated) 215
1.5.12 Ensure systemd-coredump Storage is configured (Automated) 219
1.6 Configure Command Line Warning Banners 223
1.6.1 Ensure /etc/motd is configured (Automated) 224
1.6.2 Ensure /etc/issue is configured (Automated) 226
1.6.3 Ensure /etc/issue.net is configured (Automated) 228
1.6.4 Ensure pam_motd is configured (Automated) 230
1.6.5 Ensure sshd warning Banner is configured (Automated) 233
1.6.6 Ensure access to /etc/motd is configured (Automated) 235
1.6.7 Ensure access to /etc/issue is configured (Automated) 237
1.6.8 Ensure access to /etc/issue.net is configured (Automated) 239
1.6.9 Ensure access to pam_motd file is configured (Automated) 241
1.6.10 Ensure access to sshd warning banner is configured (Automated) 243
1.6.11 Ensure update-notifier-motd.service is not in use (Automated) 245
1.6.12 Ensure update-notifier-motd.timer is not in use (Automated) 248
1.7 Configure GNOME Display Manager 251
1.7.1 Ensure GDM login banner is configured (Automated) 252
1.7.2 Ensure GDM disable-user-list is configured (Automated) 257
1.7.3 Ensure GDM screen lock is configured (Automated) 261
1.7.4 Ensure GDM automount is configured (Automated) 266
1.7.5 Ensure GDM autorun-never is configured (Automated) 270
1.7.6 Ensure XDMCP is not enabled (Automated) 273
1.7.7 Ensure Xwayland is configured (Automated) 276

2 Services 278
2.1 Configure Server Services 279
2.1.1 Ensure autofs services are not in use (Automated) 280

Internal Only - General

2.1.2 Ensure avahi daemon services are not in use (Automated) 283
2.1.3 Ensure dhcp server services are not in use (Automated) 286
2.1.4 Ensure dns server services are not in use (Automated) 288
2.1.5 Ensure dnsmasq services are not in use (Automated) 290
2.1.6 Ensure ftp server services are not in use (Automated) 292
2.1.7 Ensure ldap server services are not in use (Automated) 295
2.1.8 Ensure message access server services are not in use (Automated) 297
2.1.9 Ensure network file system services are not in use (Automated) 300
2.1.10 Ensure nis server services are not in use (Automated) 302
2.1.11 Ensure print server services are not in use (Automated) 304
2.1.12 Ensure rpcbind services are not in use (Automated) 307
2.1.13 Ensure rsync services are not in use (Automated) 310
2.1.14 Ensure samba file server services are not in use (Automated) 312
2.1.15 Ensure snmp services are not in use (Automated) 314
2.1.16 Ensure telnet server services are not in use (Automated) 317
2.1.17 Ensure tftp server services are not in use (Automated) 319
2.1.18 Ensure web proxy server services are not in use (Automated) 321
2.1.19 Ensure web server services are not in use (Automated) 323
2.1.20 Ensure xinetd services are not in use (Automated) 326
2.1.21 Ensure X window server services are not in use (Automated) 328
2.1.22 Ensure mail transfer agents are configured for local-only mode (Automated) 330
2.1.23 Ensure only approved services are listening on a network interface (Manual) 333
2.2 Configure Client Services 335
2.2.1 Ensure nis client is not installed (Automated) 336
2.2.2 Ensure rsh client is not installed (Automated) 338
2.2.3 Ensure talk client is not installed (Automated) 340
2.2.4 Ensure telnet client is not installed (Automated) 342
2.2.5 Ensure ldap client is not installed (Automated) 344
2.2.6 Ensure ftp client is not installed (Automated) 346
2.3 Configure Time Synchronization 347
2.3.1 Ensure time synchronization is in use 349
2.3.1.1 Ensure a single time synchronization daemon is in use (Automated) 350
2.3.2 Configure systemd-timesyncd 354
2.3.2.1 Ensure systemd-timesyncd configured with authorized timeserver (Automated) 356
2.3.2.2 Ensure systemd-timesyncd is enabled and running (Automated) 360
2.3.3 Configure chrony 362
2.3.3.1 Ensure chrony is configured (Automated) 363
2.3.3.2 Ensure chrony is running as user _chrony (Automated) 367
2.3.3.3 Ensure chrony is enabled and running (Automated) 369
2.4 Job Schedulers 370
2.4.1 Configure cron 371
2.4.1.1 Ensure cron daemon is enabled and active (Automated) 372
2.4.1.2 Ensure access to /etc/crontab is configured (Automated) 374
2.4.1.3 Ensure access to /etc/cron.hourly is configured (Automated) 376
2.4.1.4 Ensure access to /etc/cron.daily is configured (Automated) 378
2.4.1.5 Ensure access to /etc/cron.weekly is configured (Automated) 380
2.4.1.6 Ensure access to /etc/cron.monthly is configured (Automated) 382
2.4.1.7 Ensure access to /etc/cron.yearly is configured (Automated) 384
2.4.1.8 Ensure access to /etc/cron.d is configured (Automated) 386
2.4.1.9 Ensure access to crontab is configured (Automated) 388
2.4.2 Configure at 392
2.4.2.1 Ensure access to at is configured (Automated) 393
3 Network 395
3.1 Configure Network Devices 396
3.1.1 Ensure IPv6 status is identified (Manual) 397

Internal Only - General

3.1.2 Ensure wireless interfaces are not available (Automated) 400
3.1.3 Ensure bluetooth services are not in use (Automated) 403
3.2 Configure Network Kernel Modules 406
3.2.1 Ensure atm kernel module is not available (Automated) 407
3.2.2 Ensure can kernel module is not available (Automated) 410
3.2.3 Ensure dccp kernel module is not available (Automated) 413
3.2.4 Ensure rds kernel module is not available (Automated) 416
3.2.5 Ensure sctp kernel module is not available (Automated) 419
3.2.6 Ensure tipc kernel module is not available (Automated) 422
3.2.7 Ensure unneeded network protocol kernel modules are not available (Manual) 425
3.3 Configure Network Kernel Parameters 430
3.3.1 Configure IPv4 parameters 432
3.3.1.1 Ensure net.ipv4.ip_forward is configured (Automated) 433
3.3.1.2 Ensure net.ipv4.conf.all.forwarding is configured (Automated) 437
3.3.1.3 Ensure net.ipv4.conf.default.forwarding is configured (Automated) 441
3.3.1.4 Ensure net.ipv4.conf.all.send_redirects is configured (Automated) 445
3.3.1.5 Ensure net.ipv4.conf.default.send_redirects is configured (Automated) 449
3.3.1.6 Ensure net.ipv4.icmp_ignore_bogus_error_responses is configured (Automated) 453
3.3.1.7 Ensure net.ipv4.icmp_echo_ignore_broadcasts is configured (Automated) 457
3.3.1.8 Ensure net.ipv4.conf.all.accept_redirects is configured (Automated) 461
3.3.1.9 Ensure net.ipv4.conf.default.accept_redirects is configured (Automated) 465
3.3.1.10 Ensure net.ipv4.conf.all.secure_redirects is configured (Automated) 469
3.3.1.11 Ensure net.ipv4.conf.default.secure_redirects is configured (Automated) 473
3.3.1.12 Ensure net.ipv4.conf.all.rp_filter is configured (Automated) 477
3.3.1.13 Ensure net.ipv4.conf.default.rp_filter is configured (Automated) 481
3.3.1.14 Ensure net.ipv4.conf.all.accept_source_route is configured (Automated) 485
3.3.1.15 Ensure net.ipv4.conf.default.accept_source_route is configured (Automated) 489
3.3.1.16 Ensure net.ipv4.conf.all.log_martians is configured (Automated) 493
3.3.1.17 Ensure net.ipv4.conf.default.log_martians is configured (Automated) 497
3.3.1.18 Ensure net.ipv4.tcp_syncookies is configured (Automated) 501
3.3.1.19 Ensure net.ipv4.conf.all.route_localnet is configured (Automated) 505
3.3.2 Configure IPv6 parameters 509
3.3.2.1 Ensure net.ipv6.conf.all.forwarding is configured (Automated) 511
3.3.2.2 Ensure net.ipv6.conf.default.forwarding is configured (Automated) 515
3.3.2.3 Ensure net.ipv6.conf.all.accept_redirects is configured (Automated) 519
3.3.2.4 Ensure net.ipv6.conf.default.accept_redirects is configured (Automated) 523
3.3.2.5 Ensure net.ipv6.conf.all.accept_source_route is configured (Automated) 527
3.3.2.6 Ensure net.ipv6.conf.default.accept_source_route is configured (Automated) 531
3.3.2.7 Ensure net.ipv6.conf.all.accept_ra is configured (Automated) 535
3.3.2.8 Ensure net.ipv6.conf.default.accept_ra is configured (Automated) 539
4 Host Based Firewall 542
4.1 Configure Uncomplicated Firewall 544
4.1.1 Ensure ufw is installed (Automated) 545
4.1.2 Ensure ufw service is configured (Automated) 547
4.1.3 Ensure ufw incoming default is configured (Automated) 550
4.1.4 Ensure ufw outgoing default is configured (Automated) 552
4.1.5 Ensure ufw routed default is configured (Automated) 554
5 Access Control 555
5.1 Configure SSH Server 556
5.1.1 Ensure access to /etc/ssh/sshd_config is configured (Automated) 558
5.1.2 Ensure access to /etc/ssh/sshd_config.d/ directory is configured (Automated) 561
5.1.3 Ensure access to SSH private host key files is configured (Automated) 564
5.1.4 Ensure access to SSH public host key files is configured (Automated) 567
5.1.5 Ensure sshd access is configured (Automated) 570
5.1.6 Ensure sshd Banner is configured (Automated) 573

Internal Only - General

5.1.7 Ensure sshd Ciphers are configured (Automated) 575
5.1.8 Ensure sshd ClientAliveInterval and ClientAliveCountMax are configured (Automated) 578
5.1.9 Ensure sshd DisableForwarding is enabled (Automated) 581
5.1.10 Ensure sshd GSSAPIAuthentication is disabled (Automated) 583
5.1.11 Ensure sshd HostbasedAuthentication is disabled (Automated) 585
5.1.12 Ensure sshd IgnoreRhosts is enabled (Automated) 587
5.1.13 Ensure sshd KexAlgorithms is configured (Automated) 589
5.1.14 Ensure sshd post-quantum cryptography key exchange algorithms are configured (Manual) 592
5.1.15 Ensure sshd LoginGraceTime is configured (Automated) 595
5.1.16 Ensure sshd LogLevel is configured (Automated) 597
5.1.17 Ensure sshd MACs are configured (Automated) 599
5.1.18 Ensure sshd MaxAuthTries is configured (Automated) 602
5.1.19 Ensure sshd MaxSessions is configured (Automated) 604
5.1.20 Ensure sshd MaxStartups is configured (Automated) 606
5.1.21 Ensure sshd PermitEmptyPasswords is disabled (Automated) 608
5.1.22 Ensure sshd PermitRootLogin is disabled (Automated) 610
5.1.23 Ensure sshd PermitUserEnvironment is disabled (Automated) 612
5.1.24 Ensure sshd UsePAM is enabled (Automated) 614
5.2 Configure privilege escalation 616
5.2.1 Ensure sudo is installed (Automated) 618
5.2.2 Ensure sudo commands use pty (Automated) 621
5.2.3 Ensure sudo log file exists (Automated) 624
5.2.4 Ensure users must provide password for escalation (Automated) 627
5.2.5 Ensure re-authentication for privilege escalation is not disabled globally (Automated) 629
5.2.6 Ensure sudo timestamp_timeout is configured (Automated) 631
5.2.7 Ensure access to the su command is restricted (Automated) 633
5.3 Pluggable Authentication Modules 634
5.3.1 Configure PAM software packages 635
5.3.1.1 Ensure latest version of pam is installed (Automated) 636
5.3.1.2 Ensure latest version of libpam-modules is installed (Automated) 638
5.3.1.3 Ensure latest version of libpam-pwquality is installed (Automated) 640
5.3.1.4 Ensure latest version of cracklib-runtime is installed (Automated) 642
5.3.2 Configure pam-auth-update profiles 644
5.3.2.1 Ensure pam_unix module is enabled (Automated) 645
5.3.2.2 Ensure pam_faillock module is enabled (Automated) 647
5.3.2.3 Ensure pam_pwquality module is enabled (Automated) 650
5.3.2.4 Ensure pam_pwhistory module is enabled (Automated) 653
5.3.3 Configure PAM Arguments 655
5.3.3.1 Configure pam_faillock module 657
5.3.3.1.1 Ensure password failed attempts lockout is configured (Automated) 658
5.3.3.1.2 Ensure password unlock time is configured (Automated) 660
5.3.3.1.3 Ensure password failed attempts lockout includes root account (Automated) 663
5.3.3.2 Configure pam_pwquality module 666
5.3.3.2.1 Ensure password number of changed characters is configured (Automated) 668
5.3.3.2.2 Ensure password length is configured (Automated) 671
5.3.3.2.3 Ensure password complexity is configured (Manual) 674
5.3.3.2.4 Ensure password same consecutive characters is configured (Automated) 678
5.3.3.2.5 Ensure password maximum sequential characters is configured (Automated) 681
5.3.3.2.6 Ensure password dictionary check is enabled (Automated) 684
5.3.3.2.7 Ensure password quality checking is enforced (Automated) 686
5.3.3.2.8 Ensure password quality is enforced for the root user (Automated) 688
5.3.3.3 Configure pam_pwhistory module 691
5.3.3.3.1 Ensure password history remember is configured (Automated) 692
5.3.3.3.2 Ensure password history is enforced for the root user (Automated) 695

Internal Only - General

5.3.3.3.3 Ensure pam_pwhistory includes use_authtok (Automated) 698
5.3.3.4 Configure pam_unix module 701
5.3.3.4.1 Ensure pam_unix does not include nullok (Automated) 702
5.3.3.4.2 Ensure pam_unix does not include remember (Automated) 705
5.3.3.4.3 Ensure pam_unix includes a strong password hashing algorithm (Automated) 707
5.3.3.4.4 Ensure pam_unix includes use_authtok (Automated) 710
5.4 User Accounts and Environment 712
5.4.1 Configure shadow password suite parameters 713
5.4.1.1 Ensure password expiration is configured (Automated) 714
5.4.1.2 Ensure minimum password days is configured (Manual) 717
5.4.1.3 Ensure password expiration warning days is configured (Automated) 720
5.4.1.4 Ensure strong password hashing algorithm is configured (Automated) 722
5.4.1.5 Ensure inactive password lock is configured (Automated) 725
5.4.1.6 Ensure all users last password change date is in the past (Automated) 727
5.4.2 Configure root and system accounts and environment 729
5.4.2.1 Ensure root is the only UID 0 account (Automated) 730
5.4.2.2 Ensure root is the only GID 0 account (Automated) 732
5.4.2.3 Ensure group root is the only GID 0 group (Automated) 734
5.4.2.4 Ensure root account access is controlled (Automated) 736
5.4.2.5 Ensure root path integrity (Automated) 738
5.4.2.6 Ensure root user umask is configured (Automated) 741
5.4.2.7 Ensure system accounts do not have a valid login shell (Automated) 744
5.4.2.8 Ensure accounts without a valid login shell are locked (Automated) 746
5.4.3 Configure user default environment 748
5.4.3.1 Ensure nologin is not listed in /etc/shells (Automated) 749
5.4.3.2 Ensure default user shell timeout is configured (Automated) 751
5.4.3.3 Ensure default user umask is configured (Automated) 756
6 Logging and Auditing 760
6.1 System Logging 761
6.1.1 Configure journald 762
6.1.1.1 Configure journald 763
6.1.1.1.1 Ensure journald service is active (Automated) 764
6.1.1.1.2 Ensure systemd-journal-remote service is not in use (Automated) 766
6.1.1.1.3 Ensure journald is configured to send logs to rsyslog (Automated) 768
6.1.1.1.4 Ensure journald log file access is configured (Manual) 771
6.1.1.1.5 Ensure journald log file rotation is configured (Manual) 774
6.1.1.1.6 Ensure journald Storage is configured (Automated) 776
6.1.1.1.7 Ensure journald Compress is configured (Automated) 779
6.1.2 Configure rsyslog 783
6.1.2.1 Ensure rsyslog is installed (Automated) 784
6.1.2.2 Ensure rsyslog service is enabled and active (Automated) 786
6.1.2.3 Ensure rsyslog log file creation mode is configured (Automated) 788
6.1.2.4 Ensure rsyslog logging is configured (Manual) 790
6.1.2.5 Ensure rsyslog is configured to send logs to a remote log host (Manual) 792
6.1.2.6 Ensure rsyslog is not configured to receive logs from a remote client (Automated) 794
6.1.2.7 Ensure logrotate is configured (Manual) 797
6.1.2.8 Ensure rsyslog-gnutls is installed (Automated) 799
6.1.2.9 Ensure rsyslog forwarding uses gtls (Automated) 801
6.1.2.10 Ensure rsyslog CA certificates are configured (Manual) 803
6.1.3 Configure Logfiles 806
6.1.3.1 Ensure access to all logfiles has been configured (Automated) 807
6.2 System Auditing 812
6.2.1 Configure auditd Service 815
6.2.1.1 Ensure auditd packages are installed (Automated) 816
6.2.1.2 Ensure auditd service is enabled and active (Automated) 818

Internal Only - General

6.2.1.3 Ensure auditing for processes that start prior to auditd is enabled (Automated) 820
6.2.1.4 Ensure audit_backlog_limit is configured (Automated) 822
6.2.2 Configure Data Retention 824
6.2.2.1 Ensure audit log storage size is configured (Automated) 825
6.2.2.2 Ensure audit logs are not automatically deleted (Automated) 827
6.2.2.3 Ensure system is disabled when audit logs are full (Automated) 829
6.2.2.4 Ensure system warns when audit logs are low on space (Automated) 832
6.2.3 Configure auditd Rules 835
6.2.3.1 Ensure modification of the /etc/sudoers file is collected (Automated) 838
6.2.3.2 Ensure actions as another user are always logged (Automated) 841
6.2.3.3 Ensure events that modify the sudo log file are collected (Automated) 844
6.2.3.4 Ensure events that modify date and time information are collected (Automated) 847
6.2.3.5 Ensure events that modify sethostname and setdomainname are collected (Automated) 850
6.2.3.6 Ensure events that modify /etc/issue and /etc/issue.net are collected (Automated) .. 853
6.2.3.7 Ensure events that modify /etc/hosts and /etc/hostname are collected (Automated) 855
6.2.3.8 Ensure events that modify the system's network environment are collected (Automated) 857
6.2.3.9 Ensure events that modify /etc/NetworkManager directory are collected (Automated) 860
6.2.3.10 Ensure use of privileged commands are collected (Automated) 863
6.2.3.11 Ensure unsuccessful file access attempts are collected (Automated) 867
6.2.3.12 Ensure events that modify /etc/group information are collected (Automated) 871
6.2.3.13 Ensure events that modify /etc/passwd information are collected (Automated) 874
6.2.3.14 Ensure events that modify /etc/shadow and /etc/gshadow are collected (Automated) 877
6.2.3.15 Ensure events that modify /etc/security/opasswd are collected (Automated) 880
6.2.3.16 Ensure events that modify /etc/nsswitch.conf file are collected (Automated) 883
6.2.3.17 Ensure events that modify /etc/pam.conf and /etc/pam.d/ information are collected (Automated) 886
6.2.3.18 Ensure discretionary access control permission modification events chmod,fchmod,fchmodat,fchmodat2 are collected (Automated) 889
6.2.3.19 Ensure discretionary access control permission modification events chown,fchown,lchown,fchownat are collected (Automated) 892
6.2.3.20 Ensure discretionary access control permission modification events setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr collected (Automated) 896
6.2.3.21 Ensure successful file system mounts are collected (Automated) 900
6.2.3.22 Ensure session initiation information is collected (Automated) 903
6.2.3.23 Ensure login and logout events are collected (Automated) 906
6.2.3.24 Ensure unlink file deletion events by users are collected (Automated) 909
6.2.3.25 Ensure rename file deletion events by users are collected (Automated) 912
6.2.3.26 Ensure events that modify the system's Mandatory Access Controls are collected (Automated) 915
6.2.3.27 Ensure successful and unsuccessful attempts to use the chcon command are collected (Automated) 918
6.2.3.28 Ensure successful and unsuccessful attempts to use the setfacl command are collected (Automated) 921
6.2.3.29 Ensure successful and unsuccessful attempts to use the chacl command are collected (Automated) 924
6.2.3.30 Ensure successful and unsuccessful attempts to use the usermod command are collected (Automated) 927
6.2.3.31 Ensure kernel module loading unloading and modification is collected (Automated) 930
6.2.3.32 Ensure kernel "init_module" and "finit_module" loading unloading and modification is collected (Automated) 933

Internal Only - General

6.2.3.33 Ensure kernel "delete_module" loading unloading and modification is collected (Automated) 936
6.2.3.34 Ensure the audit configuration is loaded regardless of errors (Automated) 939
6.2.3.35 Ensure the audit configuration is immutable (Automated) 941
6.2.3.36 Ensure the running and on disk configuration is the same (Manual) 943
6.2.4 Configure auditd File Access 945
6.2.4.1 Ensure audit log files mode is configured (Automated) 946
6.2.4.2 Ensure audit log files owner is configured (Automated) 949
6.2.4.3 Ensure audit log files group owner is configured (Automated) 952
6.2.4.4 Ensure the audit log file directory mode is configured (Automated) 954
6.2.4.5 Ensure audit configuration files mode is configured (Automated) 957
6.2.4.6 Ensure audit configuration files owner is configured (Automated) 959
6.2.4.7 Ensure audit configuration files group owner is configured (Automated) 961
6.2.4.8 Ensure audit tools mode is configured (Automated) 963
6.2.4.9 Ensure audit tools owner is configured (Automated) 966
6.2.4.10 Ensure audit tools group owner is configured (Automated) 968
6.3 Configure Integrity Checking 970
6.3.1 Ensure AIDE is installed (Automated) 971
6.3.2 Ensure filesystem integrity is regularly checked (Automated) 973
6.3.3 Ensure cryptographic mechanisms are used to protect the integrity of audit tools (Automated) 975
7 System Maintenance 979
7.1 Configure system file and directory access 980
7.1.1 Ensure access to /etc/passwd is configured (Automated) 981
7.1.2 Ensure access to /etc/passwd- is configured (Automated) 983
7.1.3 Ensure access to /etc/group is configured (Automated) 985
7.1.4 Ensure access to /etc/group- is configured (Automated) 987
7.1.5 Ensure access to /etc/shadow is configured (Automated) 989
7.1.6 Ensure access to /etc/shadow- is configured (Automated) 991
7.1.7 Ensure access to /etc/gshadow is configured (Automated) 993
7.1.8 Ensure access to /etc/gshadow- is configured (Automated) 995
7.1.9 Ensure access to /etc/shells is configured (Automated) 997
7.1.10 Ensure access to /etc/security/opasswd is configured (Automated) 999
7.1.11 Ensure world writable files and directories are secured (Automated) 1001
7.1.12 Ensure no files or directories without an owner and a group exist (Automated) 1005
7.1.13 Ensure SUID and SGID files are reviewed (Manual) 1008
7.2 Local User and Group Settings 1011
7.2.1 Ensure accounts in /etc/passwd use shadowed passwords (Automated) 1012
7.2.2 Ensure /etc/shadow password fields are not empty (Automated) 1015
7.2.3 Ensure all groups in /etc/passwd exist in /etc/group (Automated) 1017
7.2.4 Ensure shadow group is empty (Automated) 1019
7.2.5 Ensure no duplicate UIDs exist (Automated) 1021
7.2.6 Ensure no duplicate GIDs exist (Automated) 1023
7.2.7 Ensure no duplicate user names exist (Automated) 1025
7.2.8 Ensure no duplicate group names exist (Automated) 1027
7.2.9 Ensure local interactive user home directories are configured (Automated) 1029
7.2.10 Ensure local interactive user dot files access is configured (Automated) 1033
7.2.11 Ensure local interactive user dot directories access is configured (Automated) 1038
Appendix: Summary Table 1045
Appendix: CIS Controls v7 IG 1 Mapped Recommendations 1070
Appendix: CIS Controls v7 IG 2 Mapped Recommendations 1076
Appendix: CIS Controls v7 IG 3 Mapped Recommendations 1089

Internal Only - General

Appendix: CIS Controls v7 Unmapped Recommendations 1102
Appendix: CIS Controls v8 IG 1 Mapped Recommendations 1103
Appendix: CIS Controls v8 IG 2 Mapped Recommendations 1110
Appendix: CIS Controls v8 IG 3 Mapped Recommendations 1123
Appendix: CIS Controls v8 Unmapped Recommendations 1136
Appendix: Change History 1137

Internal Only - General

Overview

All CIS Benchmarks™ (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 and applications for vulnerabilities and quickly updating with the latest security patches.
- End-point protection (Antivirus software, Endpoint Detection and Response (EDR), etc.).
- Logging and monitoring user and system activity.

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

Important Usage Information

All Benchmarks are available free for non-commercial use from the CIS Website. They can be used to manually assess and remediate systems and applications. In lieu of manual assessment and remediation, there are several tools available to assist with assessment:

- CIS Configuration Assessment Tool (CIS-CAT® Pro Assessor)
- CIS Benchmarks™ Certified 3rd Party Tooling

These tools make the hardening process much more scalable for large numbers of systems and applications.

NOTE: Some tooling focuses only on the Benchmark Recommendations that can be fully automated (skipping ones marked **Manual**). It is important that ***ALL*** Recommendations (**Automated** and **Manual**) be addressed since all are important for properly securing systems and are typically in scope for audits.

Key Stakeholders

Cybersecurity is a collaborative effort, and cross functional cooperation is imperative within an organization to discuss, test, and deploy Benchmarks in an effective and efficient way. The Benchmarks are developed to be best practice configuration guidelines applicable to a wide range of use cases. In some organizations, exceptions to specific Recommendations will be needed, and this team should work to prioritize the problematic Recommendations based on several factors like risk, time, cost, and labor. These exceptions should be properly categorized and documented for auditing purposes.

Internal Only - General

Apply the Correct Version of a Benchmark

Benchmarks are developed and tested for a specific set of products and versions and applying an incorrect Benchmark to a system can cause the resulting pass/fail score to be incorrect. This is due to the assessment of settings that do not apply to the target systems. To assure the correct Benchmark is being assessed:

- **Deploy the Benchmark applicable to the way settings are managed in the environment:** An example of this is the Microsoft Windows family of Benchmarks, which have separate Benchmarks for Group Policy, Intune, and Stand-alone systems based upon how system management is deployed. Applying the wrong Benchmark in this case will give invalid results.

- **Use the most recent version of a Benchmark**: This is true for all Benchmarks, but especially true for cloud technologies. Cloud technologies change frequently and using an older version of a Benchmark may have invalid methods for auditing and remediation.

Exceptions

The guidance items in the Benchmarks are called recommendations and not requirements, and exceptions to some of them are expected and acceptable. The Benchmarks strive to be a secure baseline, or starting point, for a specific technology, with known issues identified during Benchmark development are documented in the Impact section of each Recommendation. In addition, organizational, system specific requirements, or local site policy may require changes as well, or an exception to a Recommendation or group of Recommendations (e.g. A Benchmark could Recommend that a Web server not be installed on the system, but if a system's primary purpose is to function as a Webserver, there should be a documented exception to this Recommendation for that specific server).

In the end, exceptions to some Benchmark Recommendations are common and acceptable, and should be handled as follows:

- The reasons for the exception should be reviewed cross-functionally and be well documented for audit purposes.
- A plan should be developed for mitigating, or eliminating, the exception in the future, if applicable.
- If the organization decides to accept the risk of this exception (not work toward mitigation or elimination), this should be documented for audit purposes.

It is the responsibility of the organization to determine their overall security policy, and which settings are applicable to their unique needs based on the overall risk profile for the organization.

Internal Only - General

Remediation

CIS has developed Build Kits for many technologies to assist in the automation of hardening systems. Build Kits are designed to correspond to Benchmark's "Remediation" section, which provides the manual remediation steps necessary to make that Recommendation compliant to the Benchmark.

When remediating systems (changing configuration settings on deployed systems as per the Benchmark's Recommendations), please approach this with caution and test thoroughly.

The following is a reasonable remediation approach to follow:

- CIS Build Kits, or internally developed remediation methods should never be applied to production systems without proper testing.
- Proper testing consists of the following:
  - Understand the configuration (including installed applications) of the targeted systems. Various parts of the organization may need different configurations (e.g., software developers vs standard office workers).
  - Read the Impact section of the given Recommendation to help determine if there might be an issue with the targeted systems.
  - Test the configuration changes with representative lab system(s). If issues arise during testing, they can be resolved prior to deploying to any production systems.
  - When testing is complete, initially deploy to a small sub-set of production systems and monitor closely for issues. If there are issues, they can be resolved prior to deploying more broadly.
  - When the initial deployment above is completed successfully, iteratively deploy to additional systems and monitor closely for issues. Repeat this process until the full deployment is complete.

Summary

Using the Benchmarks Certified tools, working as a team with key stakeholders, being selective with exceptions, and being careful with remediation deployment, it is possible to harden large numbers of deployed systems in a cost effective, efficient, and safe manner.

NOTE: As previously stated, the PDF versions of the CIS Benchmarks™ are available for free, non-commercial use on the CIS Website. All other formats of the CIS Benchmarks™ (MS Word, Excel, and Build Kits) are available for CIS SecureSuite® members.

CIS-CAT® Pro is also available to CIS SecureSuite® members.

Internal Only - General

Target Technology Details

This document provides prescriptive guidance for establishing a secure configuration posture for **Ubuntu 26.04 LTS derived** Linux distribution systems running on x86_64 platforms.

This guide was developed and tested against **Ubuntu 26.04 LTS**

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

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

To obtain the latest version of this guide, please visit https://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 **Ubuntu 26.04 LTS derived** Linux distributions on x86_64 platforms.

Internal Only - General

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/.

Internal Only - General

Typographical Conventions

The following typographical conventions are used throughout this guide:

Convention	Meaning
`Stylized Monospace font`	Used for blocks of code, command, and script examples. Text should be interpreted exactly as presented.
`Monospace font`	Used for inline code, commands, UI/Menu selections or examples. Text should be interpreted exactly as presented.
`<Monospace font in brackets>`	Text set in angle brackets denote a variable requiring substitution for a real value.
Italic font	Used to reference other relevant settings, CIS Benchmarks and/or Benchmark Communities. Also, used to denote the title of a book, article, or other publication.
Bold font	Additional information or caveats things like **Notes**, **Warnings**, or **Cautions** (usually just the word itself and the rest of the text normal).

Internal Only - General

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.

Internal Only - General

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.

Internal Only - General

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.

Internal Only - General

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
Ron Colvin
Dave Billing
Dominic Pace
Koen Laevens
Mark Birch
Thomas Sjögren
James Trigg
Matthew Burket
Marcus Burghardt
Graham Eames
Robert McSulla
Chad Streck
Ryan Jaynes
Cory Sherman
Simon John
Steve Cobrin
Justin Brown
Mike Cross
Kleber Leal
Laurent Mesure
Andrew Box
Nathaniel Hooker

Editor
Jonathan Lewis Christopherson
Eric Pinnell
Gokhan Lus
Randie Bejar

Internal Only - General

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

The file system is generally a built-in layer used to handle the data management of the storage.

Internal Only - General

1.1.1 Configure Filesystem Kernel Modules

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

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

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

Start up scripts

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

Return values

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

Internal Only - General

1.1.1.1 Ensure cramfs kernel module is not available (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 system. If this filesystem type is not needed, disable it.

Audit:

Verify the `cramfs` kernel module is not available on the system or has been disabled.

1. Run the following script to determine if the `cramfs` kernel module is available on the system:

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

{
   l_mod_name="cramfs" l_mod_type="fs"
   while IFS= read -r l_mod_path; do
      if [ -d "$l_mod_path/${l_mod_name//-/\/}" ] &&  \
      [ -n "$(ls -A "$l_mod_path/${l_mod_name//-/\/}")" ]; then
         printf '%s\n' "$l_mod_name exists in $l_mod_path"
      fi
   done < <(readlink -e /usr/lib/modules/**/kernel/$l_mod_type \
   || readlink -e /lib/modules/**/kernel/$l_mod_type)
}
```

If nothing is returned, the `cramfs` kernel module is not available on the system and no further audit steps are required.

Note: Some systems may include the `cramfs` filesystem as part of the kernel opposed to being available as a kernel module. In this case, the above audit will not return anything. This is also considered a passing state.

If anything is returned by the above script:

2. verify the `cramfs` kernel module is not loaded and not loadable by performing the following:

Run the following command to verify the `cramfs` kernel module is not loaded:

Internal Only - General

```
# lsmod | grep 'cramfs'
```

Nothing should be returned.

Run the following command to verify the `cramfs` kernel module is not loadable:

```
# modprobe --showconfig | grep -P -- '\b(install|blacklist)\h+cramfs\b'
```

Verify the output includes:

```
blacklist cramfs
 -AND EITHER-
install cramfs /bin/false
 -OR-
install cramfs /bin/true
```

Example output:

```
blacklist cramfs
install cramfs /bin/false
```

Remediation:

Unload and disable the `cramfs` kernel module.

1. Run the following commands to unload the `cramfs` kernel module:

```
# modprobe -r cramfs 2>/dev/null
# rmmod cramfs 2>/dev/null
```

2. Perform the following to disable the `cramfs` kernel module:

Create a file ending in `.conf` with `install cramfs /bin/false` in the `/etc/modprobe.d/` directory.

Example:

```
# printf '%s\n' "" "install cramfs /bin/false" >> /etc/modprobe.d/60-
cramfs.conf
```

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

Example:

```
# printf '%s\n' "" "blacklist cramfs" >> /etc/modprobe.d/60-cramfs.conf
```

References:

1. GRID: NIX-00000000
2. NIST SP 800-53: CM-7 a
3. NIST SP 800-53A :: CM-7.1 (ii)
4. STIG ID: RHEL-08-040025 | RULE ID: SV-230498r1069314 | CAT III
5. STIG ID: OL08-00-040025 | RULE ID: SV-248833r958478 | CAT III

Internal Only - General

6. STIG ID: RHEL-09-231195 | RULE ID: SV-257880r1044951 | CAT III
7. STIG ID: ALMA-09-029940 | RULE ID: SV-269344r1050226 | CAT II
8. STIG ID: OL09-00-000045 | RULE ID: SV-271448r1091056 | CAT III
9. CCI-000381

CIS Controls:

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

Internal Only - General

1.1.1.2 Ensure freevxfs kernel module is not available (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

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

Rationale:

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

Audit:

Verify the `freevxfs` kernel module is not available on the system or has been disabled.

1. Run the following script to determine if the `freevxfs` kernel module is available on the system:

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

{
   l_mod_name="freevxfs" l_mod_type="fs"
   while IFS= read -r l_mod_path; do
      if [ -d "$l_mod_path/${l_mod_name//-/\/}" ] &&  \
      [ -n "$(ls -A "$l_mod_path/${l_mod_name//-/\/}")" ]; then
         printf '%s\n' "$l_mod_name exists in $l_mod_path"
      fi
   done < <(readlink -e /usr/lib/modules/**/kernel/$l_mod_type \
   || readlink -e /lib/modules/**/kernel/$l_mod_type)
}
```

If nothing is returned, the `freevxfs` kernel module is not available on the system and no further audit steps are required.

Note: Some systems may include the `freevxfs` filesystem as part of the kernel opposed to being available as a kernel module. In this case, the above audit will not return anything. This is also considered a passing state.

If anything is returned by the above script:

2. Verify the `freevxfs` kernel module is not loaded and not loadable by performing the following:

Internal Only - General

Run the following command to verify the `freevxfs` kernel module is not loaded:

```
# lsmod | grep 'freevxfs'
```

Nothing should be returned.

Run the following command to verify the `freevxfs` kernel module is not loadable:

```
# modprobe --showconfig | grep -P -- '\b(install|blacklist)\h+freevxfs\b'
```

Verify the output includes:

```
blacklist freevxfs
 -AND-
install freevxfs /bin/false
 -OR-
install freevxfs /bin/true
```

Example output:

```
blacklist freevxfs
install freevxfs /bin/false
```

Remediation:

Unload and disable the `freevxfs` kernel module.

1. Run the following commands to unload the `freevxfs` kernel module:

```
# modprobe -r freevxfs 2>/dev/null
# rmmod freevxfs 2>/dev/null
```

2. Perform the following to disable the `freevxfs` kernel module:

Create a file ending in `.conf` with `install freevxfs /bin/false` in the `/etc/modprobe.d/` directory.

Example:

```
# printf '%s\n' "" "install freevxfs /bin/false" >> /etc/modprobe.d/60-
freevxfs.conf
```

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

Example:

```
# printf '%s\n' "" "blacklist freevxfs" >> /etc/modprobe.d/60-freevxfs.conf
```

References:

1. GRID: NIX-00000001
2. NIST SP 800-53 :: CM-7
3. NIST SP 800-53A :: CM-7.1 (ii)
4. CCI-000381

Internal Only - General

CIS Controls:

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

Internal Only - General

1.1.1.3 Ensure hfs kernel module is not available (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

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

Rationale:

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

Audit:

Verify the `hfs` kernel module is not available on the system or has been disabled.

1. Run the following script to determine if the `hfs` kernel module is available on the system:

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

{
   l_mod_name="hfs" l_mod_type="fs"
   while IFS= read -r l_mod_path; do
      if [ -d "$l_mod_path/${l_mod_name//-/\/}" ] &&  \
      [ -n "$(ls -A "$l_mod_path/${l_mod_name//-/\/}")" ]; then
         printf '%s\n' "$l_mod_name exists in $l_mod_path"
      fi
   done < <(readlink -e /usr/lib/modules/**/kernel/$l_mod_type \
   || readlink -e /lib/modules/**/kernel/$l_mod_type)
}
```

If nothing is returned, the `hfs` kernel module is not available on the system and no further audit steps are required.

Note: Some systems may include the `hfs` filesystem as part of the kernel opposed to being available as a kernel module. In this case, the above audit will not return anything. This is also considered a passing state.

If anything is returned by the above script:

2. Verify the `hfs` kernel module is not loaded and not loadable by performing the following:

Run the following command to verify the `hfs` kernel module is not loaded:

Internal Only - General

```
# lsmod | grep 'hfs'
```

Nothing should be returned.

Run the following command to verify the `hfs` kernel module is not loadable:

```
# modprobe --showconfig | grep -P -- '\b(install|blacklist)\h+hfs\b'
```

Verify the output includes:

```
blacklist hfs
 -AND-
install hfs /bin/false
 -OR-
install hfs /bin/true
```

Example output:

```
blacklist hfs
install hfs /bin/false
```

Remediation:

Unload and disable the `hfs` kernel module.

1. Run the following commands to unload the `hfs` kernel module:

```
# modprobe -r hfs 2>/dev/null
# rmmod hfs 2>/dev/null
```

2. Perform the following to disable the `hfs` kernel module:

Create a file ending in `.conf` with `install hfs /bin/false` in the `/etc/modprobe.d/` directory.

Example:

```
# printf '%s\n' "" "install hfs /bin/false" >> /etc/modprobe.d/60-hfs.conf
```

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

Example:

```
# printf '%s\n' "" "blacklist hfs" >> /etc/modprobe.d/60-hfs.conf
```

References:

1. GRID: NIX-00000002
2. NIST SP 800-53 :: CM-7
3. NIST SP 800-53A :: CM-7.1 (ii)
4. CCI-000381

Internal Only - General

CIS Controls:

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

Internal Only - General

1.1.1.4 Ensure hfsplus kernel module is not available (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

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

Rationale:

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

Audit:

Verify the `hfsplus` kernel module is not available on the system or has been disabled.

1. Run the following script to determine if the `hfsplus` kernel module is available on the system:

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

{
   l_mod_name="hfsplus" l_mod_type="fs"
   while IFS= read -r l_mod_path; do
      if [ -d "$l_mod_path/${l_mod_name//-/\/}" ] &&  \
      [ -n "$(ls -A "$l_mod_path/${l_mod_name//-/\/}")" ]; then
         printf '%s\n' "$l_mod_name exists in $l_mod_path"
      fi
   done < <(readlink -e /usr/lib/modules/**/kernel/$l_mod_type \
   || readlink -e /lib/modules/**/kernel/$l_mod_type)
}
```

If nothing is returned, the `hfsplus` kernel module is not available on the system and no further audit steps are required.

Note: Some systems may include the `hfsplus` filesystem as part of the kernel opposed to being available as a kernel module. In this case, the above audit will not return anything. This is also considered a passing state.

If anything is returned by the above script:

2. Verify the `hfsplus` kernel module is not loaded and not loadable by performing the following:

Run the following command to verify the `hfsplus` kernel module is not loaded:

Internal Only - General

```
# lsmod | grep 'hfsplus'
```

Nothing should be returned.

Run the following command to verify the `hfsplus` kernel module is not loadable:

```
# modprobe --showconfig | grep -P -- '\b(install|blacklist)\h+hfsplus\b'
```

Verify the output includes:

```
blacklist hfsplus
 -AND-
install hfsplus /bin/false
 -OR-
install hfsplus /bin/true
```

Example output:

```
blacklist hfsplus
install hfsplus /bin/false
```

Remediation:

Unload and disable the `hfsplus` kernel module.

1. Run the following commands to unload the `hfsplus` kernel module:

```
# modprobe -r hfsplus 2>/dev/null
# rmmod hfsplus 2>/dev/null
```

2. Perform the following to disable the `hfsplus` kernel module:

Create a file ending in `.conf` with `install hfsplus /bin/false` in the `/etc/modprobe.d/` directory.

Example:

```
# printf '%s\n' "" "install hfsplus /bin/false" >> /etc/modprobe.d/60-
hfsplus.conf
```

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

Example:

```
# printf '%s\n' "" "blacklist hfsplus" >> /etc/modprobe.d/60-hfsplus.conf
```

References:

1. GRID: NIX-00000003
2. NIST SP 800-53 :: CM-7
3. NIST SP 800-53A :: CM-7.1 (ii)
4. CCI-000381

Internal Only - General

CIS Controls:

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

Internal Only - General

1.1.1.5 Ensure jffs2 kernel module is not available (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

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

Rationale:

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

Audit:

Verify the `jffs2` kernel module is not available on the system or has been disabled.

1. Run the following script to determine if the `jffs2` kernel module is available on the system:

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

{
   l_mod_name="jffs2" l_mod_type="fs"
   while IFS= read -r l_mod_path; do
      if [ -d "$l_mod_path/${l_mod_name//-/\/}" ] &&  \
      [ -n "$(ls -A "$l_mod_path/${l_mod_name//-/\/}")" ]; then
         printf '%s\n' "$l_mod_name exists in $l_mod_path"
      fi
   done < <(readlink -e /usr/lib/modules/**/kernel/$l_mod_type \
   || readlink -e /lib/modules/**/kernel/$l_mod_type)
}
```

If nothing is returned, the `jffs2` kernel module is not available on the system and no further audit steps are required.

Note: Some systems may include the `jffs2` filesystem as part of the kernel opposed to being available as a kernel module. In this case, the above audit will not return anything. This is also considered a passing state.

If anything is returned by the above script:

2. Verify the `jffs2` kernel module is not loaded and not loadable by performing the following:

Run the following command to verify the `jffs2` kernel module is not loaded:

Internal Only - General

```
lsmod | grep 'jffs2'
```

Nothing should be returned.

Run the following command to verify the `jffs2` kernel module is not loadable:

```
modprobe --showconfig | grep -P -- '\b(install|blacklist)\h+jffs2\b'
```

Verify the output includes:

```
blacklist jffs2
 -AND-
install jffs2 /bin/false
 -OR-
install jffs2 /bin/true
```

Example output:

```
blacklist jffs2
install jffs2 /bin/false
```

Remediation:

Unload and disable the `jffs2` kernel module.

1. Run the following commands to unload the `jffs2` kernel module:

```
# modprobe -r jffs2 2>/dev/null
# rmmod jffs2 2>/dev/null
```

2. Perform the following to disable the `jffs2` kernel module:

Create a file ending in `.conf` with `install jffs2 /bin/false` in the `/etc/modprobe.d/` directory.

Example:

```
# printf '%s\n' "" "install jffs2 /bin/false" >> /etc/modprobe.d/60-
jffs2.conf
```

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

Example:

```
# printf '%s\n' "" "blacklist jffs2" >> /etc/modprobe.d/60-jffs2.conf
```

References:

1. GRID: NIX-00000004
2. NIST SP 800-53 :: CM-7
3. NIST SP 800-53A :: CM-7.1 (ii)
4. CCI-000381

Internal Only - General

CIS Controls:

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

Internal Only - General

1.1.1.6 Ensure overlay kernel module is not available (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

The Linux `overlay` kernel module, commonly referred to as Overlayfs, is a union mount filesystem. It allows the user to overlay one file system on top of another, and creates a single, unified view.

Rationale:

The `overlay` kernel module has known CVE's: CVE-2023-32629, CVE-2023-2640, and CVE-2023-0386. Disabling the `overlay` kernel module reduces the local attack surface by removing support for unnecessary filesystem types and mitigates potential risks associated with unauthorized execution of setuid files, enhancing the overall system security.

Impact:

WARNING: If Container applications such as Docker, Kubernetes, Podman, Linux Containers (LXC), etc. are in use proceed with caution and consider the impact on containerized workloads, as disabling the `overlay` may severely disrupt containerization.

Audit:

Verify the `overlay` kernel module is not available on the system or has been disabled.

1. Run the following script to determine if the `overlay` kernel module is available on the system:

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

{
   l_mod_name="overlayfs" l_mod_type="fs"
   while IFS= read -r l_mod_path; do
      if [ -d "$l_mod_path/${l_mod_name//-/\/}" ] &&  \
      [ -n "$(ls -A "$l_mod_path/${l_mod_name//-/\/}")" ]; then
         printf '%s\n' "$l_mod_name exists in $l_mod_path"
      fi
   done < <(readlink -e /usr/lib/modules/**/kernel/$l_mod_type \
   || readlink -e /lib/modules/**/kernel/$l_mod_type)
}
```

If nothing is returned, the `overlay` kernel module is not available on the system and no further audit steps are required.

Internal Only - General

Note: Some systems may include the `overlay` kernel module as part of the kernel opposed to being available as a kernel module. In this case, the above audit will not return anything. This is also considered a passing state.

If anything is returned by the above script:

2. Verify the `overlay` kernel module is not loaded and not loadable by performing the following:

Run the following command to verify the `overlay` kernel module is not loaded:

```
# lsmod | grep 'overlay'
```

Nothing should be returned

Run the following command to verify the `overlay` kernel module is not loadable:

```
# modprobe --showconfig | grep -P -- '\b(install|blacklist)\h+overlay\b'
```

Verify the output includes:

```
blacklist overlay
 -AND-
install overlay /bin/false
 -OR-
install overlay /bin/true
```

Example output:

```
blacklist overlay
install overlay /bin/false
```

Remediation:

Unload and disable the `overlay` kernel module.

1. Run the following commands to unload the `overlay` kernel module:

```
# modprobe -r overlay 2>/dev/null
# rmmod overlay 2>/dev/null
```

2. Perform the following to disable the `overlay` kernel module:

Create a file ending in `.conf` with `install overlay /bin/false` in the `/etc/modprobe.d/` directory

Example:

```
# printf '%s\n' "" "install overlay /bin/false" >> /etc/modprobe.d/60-
overlay.conf
```

Create a file ending in `.conf` with `blacklist overlay` in the `/etc/modprobe.d/` directory

Internal Only - General

Example:

```
# printf '%s\n' "" "blacklist overlay" >> /etc/modprobe.d/60-overlay.conf
```

References:

1. GRID: NIX-00000005
2. NIST SP 800-53 :: CM-7
3. NIST SP 800-53A :: CM-7.1 (ii)
4. https://docs.kernel.org/filesystems/overlayfs.html
5. https://wiki.archlinux.org/title/Overlay_filesystem
6. https://cve.mitre.org/cgi-bin/cvekey.cgi?keyword=overlayfs
7. CCI-000381

CIS Controls:

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

Internal Only - General

1.1.1.7 Ensure squashfs kernel module is not available (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

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

Rationale:

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

Impact:

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

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

Audit:

Verify the `squashfs` kernel module is not available on the system or has been disabled.

1. Run the following script to determine if the `squashfs` kernel module is available on the system:

Internal Only - General

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

{
   l_mod_name="squashfs" l_mod_type="fs"
   while IFS= read -r l_mod_path; do
      if [ -d "$l_mod_path/${l_mod_name//-/\/}" ] &&  \
      [ -n "$(ls -A "$l_mod_path/${l_mod_name//-/\/}")" ]; then
         printf '%s\n' "$l_mod_name exists in $l_mod_path"
      fi
   done < <(readlink -e /usr/lib/modules/**/kernel/$l_mod_type \
   || readlink -e /lib/modules/**/kernel/$l_mod_type)
}
```

If nothing is returned, the squashfs kernel module is not available on the system and no further audit steps are required.

Note: Some systems may include the squashfs filesystem as part of the kernel opposed to being available as a kernel module. In this case, the above audit will not return anything. This is also considered a passing state.

If anything is returned by the above script:

2. Verify the squashfs kernel module is not loaded and not loadable by performing the following:

Run the following command to verify the squashfs kernel module is not loaded:

```
# lsmod | grep 'squashfs'
```

Nothing should be returned.

Run the following command to verify the squashfs kernel module is not loadable:

```
# modprobe --showconfig | grep -P -- '\b(install|blacklist)\h+squashfs\b'
```

Verify the output includes:

```
blacklist squashfs
 -AND-
install squashfs /bin/false
 -OR-
install squashfs /bin/true
```

Example output:

```
blacklist squashfs
install squashfs /bin/false
```

Remediation:

Unload and disable the squashfs kernel module.

1. Run the following commands to unload the squashfs kernel module:

Internal Only - General

```
# modprobe -r squashfs 2>/dev/null
# rmmod squashfs 2>/dev/null
```

2. Perform the following to disable the `squashfs` kernel module:

Create a file ending in `.conf` with `install squashfs /bin/false` in the `/etc/modprobe.d/` directory.

Example:

```
# printf '%s\n' "" "install squashfs /bin/false" >> /etc/modprobe.d/60-squashfs.conf
```

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

Example:

```
# printf '%s\n' "" "blacklist squashfs" >> /etc/modprobe.d/60-squashfs.conf
```

References:

1. GRID: NIX-00000006
2. NIST SP 800-53 :: CM-7
3. NIST SP 800-53A :: CM-7.1 (ii)
4. STIG ID: ALMA-09-030160 | RULE ID: SV-269346r1050228 | CAT II
5. CCI-000381

CIS Controls:

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

Internal Only - General

1.1.1.8 Ensure udf kernel module is not available (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - 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.

Impact:

Microsoft Azure requires the usage of udf.

udf **should not** be disabled on systems run on Microsoft Azure.

Audit:

Verify the udf kernel module is not available on the system or has been disabled.

1. Run the following script to determine if the udf kernel module is available on the system:

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

{
   l_mod_name="udf" l_mod_type="fs"
   while IFS= read -r l_mod_path; do
      if [ -d "$l_mod_path/${l_mod_name//-/\/}" ] &&  \
      [ -n "$(ls -A "$l_mod_path/${l_mod_name//-/\/}")" ]; then
         printf '%s\n' "$l_mod_name exists in $l_mod_path"
      fi
   done < <(readlink -e /usr/lib/modules/**/kernel/$l_mod_type \
   || readlink -e /lib/modules/**/kernel/$l_mod_type)
}
```

If nothing is returned, the udf kernel module is not available on the system and no further audit steps are required.

Note: Some systems may include the udf filesystem as part of the kernel opposed to being available as a kernel module. In this case, the above audit will not return anything. This is also considered a passing state.

Internal Only - General

If anything is returned by the above script:

2. Verify the udf kernel module is not loaded and not loadable by performing the following:

Run the following command to verify the udf kernel module is not loaded:

```
# lsmod | grep 'udf'
```

Nothing should be returned

Run the following command to verify the udf kernel module is not loadable:

```
# modprobe --showconfig | grep -P -- '\b(install|blacklist)\h+udf\b'
```

Verify the output includes:

```
blacklist udf
 -AND-
install udf /bin/false
 -OR-
install udf /bin/true
```

Example output:

```
blacklist udf
install udf /bin/false
```

Remediation:

Unload and disable the udf kernel module.

1. Run the following commands to unload the udf kernel module:

```
# modprobe -r udf 2>/dev/null
# rmmod udf 2>/dev/null
```

2. Perform the following to disable the udf kernel module:

Create a file ending in .conf with install udf /bin/false in the /etc/modprobe.d/ directory

Example:

```
# printf '%s\n' "" "install udf /bin/false" >> /etc/modprobe.d/60-udf.conf
```

Create a file ending in .conf with blacklist udf in the /etc/modprobe.d/ directory

Example:

Internal Only - General

```
# printf '%s\n' "" "blacklist udf" >> /etc/modprobe.d/60-udf.conf
```

References:

1. GRID: NIX-00000007
2. NIST SP 800-53 :: CM-7
3. NIST SP 800-53A :: CM-7.1 (ii)
4. STIG ID: ALMA-09-030380 | RULE ID: SV-269348r1050230 | CAT II
5. CCI-000381

CIS Controls:

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

Internal Only - General

1.1.1.9 Ensure firewire-core kernel module is not available (Automated)

Profile Applicability:

- Level 1 - Server
- Level 2 - Workstation

Description:

The IEEE 1394 (FireWire) is a serial bus standard for high-speed real-time communication.

Rationale:

Disabling FireWire protects the system against exploitation of any flaws in its implementation.

Audit:

Verify the `firewire-core` kernel module is not available on the system or has been disabled.

1. Run the following script to determine if the `firewire-core` kernel module is available on the system:

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

{
   l_mod_name="firewire" l_mod_type="drivers"
   while IFS= read -r l_mod_path; do
      if [ -d "$l_mod_path/${l_mod_name//-/\/}" ] &&  \
      [ -n "$(ls -A "$l_mod_path/${l_mod_name//-/\/}")" ]; then
         printf '%s\n' "$l_mod_name exists in $l_mod_path"
      fi
   done < <(readlink -e /usr/lib/modules/**/kernel/$l_mod_type \
   || readlink -e /lib/modules/**/kernel/$l_mod_type)
}
```

If nothing is returned, the `firewire-core` kernel module is not available on the system and no further audit steps are required.

Note: Some systems may include the `firewire-core` filesystem as part of the kernel opposed to being available as a kernel module. In this case, the above audit will not return anything. This is also considered a passing state.

If anything is returned by the above script:

2. Verify the `firewire-core` kernel module is not loaded and not loadable by performing the following:

Internal Only - General

Run the following command to verify the `firewire-core` kernel module is not loaded:

```
# lsmod | grep -P -- 'firewire(_|-)core'
```

Nothing should be returned

Run the following command to verify the `firewire-core` kernel module is not loadable:

```
# modprobe --showconfig | grep -P -- '\b(install|blacklist)\h+firewire(_|-
)core\b'
```

Verify the output includes:

```
blacklist firewire_core
 -AND-
install firewire_core /bin/false
 -OR-
install firewire_core /bin/true
```

Example output:

```
blacklist firewire_core
install firewire_core /bin/false
```

Remediation:

Unload and disable the `firewire-core` kernel module.

1. Run the following commands to unload the `firewire-core` kernel module:

```
# modprobe -r firewire-core 2>/dev/null
# rmmod firewire-core 2>/dev/null
```

2. Perform the following to disable the `firewire-core` kernel module:

Create a file ending in `.conf` with `install firewire-core /bin/false` in the `/etc/modprobe.d/` directory

Example:

```
# printf '%s\n' "" "install firewire-core /bin/false" >> /etc/modprobe.d/60-
firewire-core.conf
```

Create a file ending in `.conf` with `blacklist firewire-core` in the `/etc/modprobe.d/` directory

Example:

```
# printf '%s\n' "" "blacklist firewire-core" >> /etc/modprobe.d/60-firewire-
core.conf
```

References:

1. GRID: NIX-00000008
2. NIST SP 800-53 :: CM-7

Internal Only - General

3. NIST SP 800-53A :: CM-7.1 (ii)
4. STIG ID: RHEL-08-040026 | RULE ID: SV-230499r1069315 | CAT III
5. STIG ID: OL08-00-040026 | RULE ID: SV-248834r958478 | CAT III
6. STIG ID: RHEL-09-213055 | RULE ID: SV-257806r1044859 | CAT II
7. STIG ID: OL09-00-000042 | RULE ID: SV-271445r1091047 | CAT II
8. CCI-000381

CIS Controls:

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

Internal Only - General

1.1.1.10 Ensure usb-storage kernel module is not available (Automated)

Profile Applicability:

- Level 1 - Server
- Level 2 - Workstation

Description:

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

Rationale:

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

Satisfies: SRG-OS-000114-GPOS-00059, SRG-OS-000378-GPOS-00163

Impact:

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

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

Audit:

Verify the usb-storage kernel module is not available on the system or has been disabled.

Run the following script to determine if the usb-storage kernel module is available on the system: #!/usr/bin/env bash

{ l_mod_name="usb-storage" l_mod_type="drivers" while IFS= read -r l_mod_path; do if [-d "$l_mod_path/${l_mod_name//-//}"] &&
[-n "$(ls -A "$l_mod_path/${l_mod_name//-//}")"]; then printf '%s\n' "$l_mod_name exists in $l_mod_path" fi done < <(readlink -e /usr/lib/modules/**/kernel/$l_mod_type || readlink -e /lib/modules/**/kernel/$l_mod_type) }

If nothing is returned, the usb-storage kernel module is not available on the system and no further audit steps are required.

Internal Only - General

Note: Some systems may include the usb-storage filesystem as part of the kernel opposed to being available as a kernel module. In this case, the above audit will not return anything. This is also considered a passing state.

If anything is returned by the above script:

Verify the usb-storage kernel module is not loaded and not loadable by performing the following:

Run the following command to verify the usb-storage kernel module is not loaded:

lsmod | grep -P -- 'usb(_|-)storage'

Nothing should be returned

Run the following command to verify the usb-storage kernel module is not loadable:

modprobe --showconfig | grep -P -- '\b(install|blacklist)\h+usb(_|-)storage\b'

Verify the output includes:

blacklist usb_storage -AND- install usb_storage /bin/false -OR- install usb_storage /bin/true

Example output:

blacklist usb_storage install usb_storage /bin/false

Remediation:

Unload and disable the `usb-storage` kernel module.

1. Run the following commands to unload the `usb-storage` kernel module:

```
# modprobe -r usb-storage 2>/dev/null
# rmmod usb-storage 2>/dev/null
```

2. Perform the following to disable the `usb-storage` kernel module:

Create a file ending in `.conf` with `install usb_storage /bin/false` in the `/etc/modprobe.d/` directory

Example:

```
# printf '%s\n' "" "install usb_storage /bin/false" >> /etc/modprobe.d/60-usb-storage.conf
```

Create a file ending in `.conf` with `blacklist usb_storage` in the `/etc/modprobe.d/` directory

Internal Only - General

Example:

```
# printf '%s\n' "" "blacklist usb_storage" >> /etc/modprobe.d/60-usb-
storage.conf
```

References:

1. GRID: NIX-00000009
2. NIST SP 800-53 :: IA-3 | NIST SP 800-53 :: CM-7 (9) (b)
3. NIST SP 800-53A :: IA-3.1 (ii)
4. STIG ID: RHEL-08-040080 | RULE ID: SV-230503r1069316 | CAT II
5. STIG ID: OL08-00-040080 | RULE ID: SV-248837r986386 | CAT II
6. STIG ID: RHEL-09-291010 | RULE ID: SV-258034r1051267 | CAT II
7. STIG ID: ALMA-09-031370 | RULE ID: SV-269357r1050240 | CAT II
8. STIG ID: OL09-00-000047 | RULE ID: SV-271450r1092466 | CAT II
9. STIG ID: OL09-00-002332 | RULE ID: SV-271702r1091818 | CAT II
10. STIG ID: UBTU-22-291010 | RULE ID: SV-260540r986276 | CAT II
11. STIG ID: UBTU-24-300039 | RULE ID: SV-270718r1067128 | CAT II
12. STIG ID: SLES-15-010480 | RULE ID: SV-234856r958820 | CAT II
13. CCI-000778 | CCI-001958 | CCI-003959

Additional Information:

An alternative solution to disabling the usb-storage module may be found in USBGuard.

Use of USBGuard and construction of USB device policies should be done in alignment with site policy.

CIS Controls:

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

Internal Only - General

1.1.1.11 Ensure unused filesystems kernel modules are not available (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Filesystem kernel modules are pieces of code that can be dynamically loaded into the Linux kernel to extend its filesystem capabilities, or so-called base kernel, of an operating system. Filesystem kernel modules are typically used to add support for new hardware (as device drivers), or for adding system calls.

Rationale:

While loadable filesystem kernel modules are a convenient method of modifying the running kernel, this can be abused by attackers on a compromised system to prevent detection of their processes or files, allowing them to maintain control over the system. Many rootkits make use of loadable filesystem kernel modules in this way.

Removing support for unneeded filesystem types reduces the local attack surface of the system. If this filesystem type is not needed, disable it. The following filesystem kernel modules have known CVE's and should be made unavailable if no dependencies exist:

- `ceph` - CVE-2022-0670
- `cifs` - CVE-2022-29869
- `exfat` CVE-2022-29973
- `ext` CVE-2022-1184
- `fat` CVE-2022-22043
- `fscache` CVE-2022-3630
- `fuse` CVE-2023-0386
- `gfs2` CVE-2023-3212
- `nfs_common` CVE-2023-6660
- `nfsd` CVE-2022-43945
- `smbfs_common` CVE-2022-2585

Impact:

This list may be quite extensive and covering all edges cases is difficult. Therefore, it's crucial to carefully consider the implications and dependencies before making any changes to the filesystem kernel module configurations.

Internal Only - General

Audit:

Verify filesystem kernel modules that are not in use on the system are disabled. Review currently loaded kernel modules to ensure they are needed and conform to local site policy.

Note: The "Audit script" included at the bottom of the "Audit Procedure" will perform the steps included below.

1. Run the following command to generate a list of filesystem types currently mounted on the system.

```
# findmnt -Dkerno fstype | sort -u
```

Review the output. These filesystem types should either be excluded from the audit or if no longer needed, unmounted and disabled.

Example output:

```
devtmpfs
ext4
tmpfs
vfat
```

2. Run the following script to generate a list of all file system kernel modules currently installed on the system:

Internal Only - General

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

{
   l_search="$(readlink -e /usr/lib/modules/ || readlink -e /lib/modules/)"
   a_module=()
   while IFS= read -r -d $'\0' l_module_dir; do
      if [ ! "$(basename "$l_module_dir")" = "nls" ]; then
         while IFS= read -r -d $'\0' l_module_file; do
            l_mname="$(basename "$l_module_file" | cut -d'.' -f1)"
            if [ -f "$l_module_file" ] && ! grep -Psiq -- '\b'"$l_mname"'\b'
\
            <<< "${a_module[*]}"; then
               a_module+=("$l_mname")
            fi
         done < <(find -L "$l_module_dir" -mindepth 1 -maxdepth 1 -type f -
print0)
      fi
   done < <(find "$l_search"/**/kernel/fs/ -mindepth 1 -maxdepth 1 -type d !
-empty -print0)
   printf '%s\n' "" "${a_module[@]}"
}
```

The output is a list of the kernel modules included in the `/usr/lib/modules/{KERNEL_VERSION}/kernel/fs/` or `/lib/modules/{KERNEL_VERSION}/kernel/fs/` directories. This list has been sorted to be unique values, ignore items in the `nls` directory, and added to an array called `a_module`.

3. Run the following command to generate a list of modules currently loaded into the kernel:

```
# lsmod
```

Note: This command lists all loaded kernel modules, not just filesystem modules. When comparing against the step 2 output, focus only on modules that also appear in the step 2 list (filesystem kernel modules under kernel/fs/).

4. compare the output from the list generated in step 2 to the list of loaded modules generated by the `lsmod` command in step 3. Any module appearing in both lists is currently loaded into the kernel. Review these to verify they are needed and approved by local site policy
5. Run the following command to generate a list of modules that are set to be Deny Listed: `blacklist {MODULE_NAME}`, and/or set be be not loadable: `install {MODULE_NAME} /bin/false` or `install {MODULE_NAME} /bin/true`:

```
# modprobe --showconfig | grep -Pi -- '^\h*(blacklist|install)\h+'
```

6. Verify all modules in the output from the list generated in step 2 are both Deny Listed: `blacklist {MODULE_NAME}` and set to be not loadable: `install {MODULE_NAME} /bin/false` or `install {MODULE_NAME} /bin/true`.

Internal Only - General

Audit Script:

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

{
   a_check=() a_output=() a_module=() a_output2=() a_output3=()
   l_search="$(readlink -e /usr/lib/modules/ || readlink -e /lib/modules/)"
   IFS=$'\n' read -r -d '' -a a_mounted < <(findmnt -Dkerno fstype \
   | sort -u && printf '\0' )
   IFS=$'\n' read -r -d '' -a a_lsmod < <(lsmod | awk '{print $1}' && printf
'\0' )
   IFS=$'\n' read -r -d '' -a a_showconfig < <(modprobe --showconfig | \
   grep -Pi -- '^\h*(blacklist|install)\h+' && printf '\0')
   while IFS= read -r -d $'\0' l_module_dir; do
      if [ ! "$(basename "$l_module_dir")" = "nls" ]; then
         while IFS= read -r -d $'\0' l_module_file; do
            l_mname="$(basename "$l_module_file" | cut -d'.' -f1)"
            if [ -f "$l_module_file" ] && ! grep -Psiq -- '\b'"$l_mname"'\b'
<<< "${a_module[*]}"; then
               a_module+=("$l_mname")
            fi
         done < <(find -L "$l_module_dir" -mindepth 1 -maxdepth 1 -type f -
print0)
      fi
   done < <(find "$l_search"/**/kernel/fs/ -mindepth 1 -maxdepth 1 -type d !
-empty -print0)
   for l_module in "${a_module[@]}"; do
      if grep -Psoiq -- '\b'"$l_module"'\b' <<< "${a_mounted[*]}"; then
         a_output+=("  - \"$l_module\"")
      elif grep -Psoiq -- '\b'"$l_module"'\b' <<< "${a_lsmod[*]}"; then
         a_output2+=("  - \"$l_module\"")
      elif ! grep -Psioq -- '\binstall\h+'"${l_module//-/_}"'\h+\H+\b' <<<
"${a_showconfig[*]}" || \
      ! grep -Psioq -- '\bblacklist\h+'"${l_module//-/_}"'\b' <<<
"${a_showconfig[*]}"; then
         a_output3+=("  - \"$l_module\"")
      fi
   done
   [ "${#a_output[@]}" -gt 0 ] && printf '%s\n' "" \
   "- There are \"${#a_output[@]}\" kernel modules currently mounted:"
"${a_output[@]}"
   [ "${#a_output2[@]}" -gt 0 ] && printf '%s\n' "" \
   "- There are \"${#a_output2[@]}\" kernel modules currently loaded:"
"${a_output2[@]}"
   [ "${#a_output3[@]}" -gt 0 ] && printf '%s\n' "" \
   "- There are \"${#a_output3[@]}\" Kernel modules currently loadable:"
"${a_output3[@]}"
}
```

WARNING: disabling or denylisting filesystem modules that are in use on the system may be FATAL. It is extremely important to thoroughly review this list.

Remediation:

- IF - the module is available in the running kernel:

Internal Only - General

- Unload the filesystem kernel module from the kernel
- Create a file ending in `.conf` with install filesystem kernel modules `/bin/false` in the `/etc/modprobe.d/` directory
- Create a file ending in `.conf` with deny list filesystem kernel modules in the `/etc/modprobe.d/` directory

WARNING: unloading, disabling or denylisting filesystem modules that are in use on the system may be FATAL. It is extremely important to thoroughly review the filesystems returned by the audit before following the remediation procedure.

Example of unloading the `gfs2`kernel module:

```
# modprobe -r gfs2 2>/dev/null
# rmmod gfs2 2>/dev/null
```

Example of fully disabling the `gfs2` kernel module:

```
# printf '%s\n' "" "blacklist gfs2" "install gfs2 /bin/false" >>
/etc/modprobe.d/gfs2.conf
```

Note:

- Disable a kernel module by modifying the command above for each unused filesystem kernel module
- The example `gfs2` must be updated with the appropriate module name for the command or example script bellow to run correctly.

Below is an example script that can be modified to use on various filesystem kernel modules manual remediation process:

Example Script

Internal Only - General

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

{
   a_output2=(); a_output3=(); l_dl="" # Initialize arrays and clear
variables
   l_mod_name="gfs2" # set module name
   l_mod_type="fs" # set module type
   l_mod_path="$(readlink -f /usr/lib/modules/**/kernel/$l_mod_type ||
readlink -f /lib/modules/**/kernel/$l_mod_type)"
   f_module_fix()
   {
      l_dl="y" # Set to ignore duplicate checks
      a_showconfig=() # Create array with modprobe output
      while IFS= read -r l_showconfig; do
         a_showconfig+=("$l_showconfig")
      done < <(modprobe --showconfig | grep -P --
'\b(install|blacklist)\h+'"${l_mod_name//-/_}"'\b')
      if  lsmod | grep "$l_mod_name" &> /dev/null; then # Check if the module
is currently loaded
         a_output2+=(" - unloading kernel module: \"$l_mod_name\"")
         modprobe -r "$l_mod_name" 2>/dev/null; rmmod "$l_mod_name"
2>/dev/null
      fi
      if ! grep -Pq -- '\binstall\h+'"${l_mod_name//-
/_}"'\h+(\/usr)?\/bin\/(true|false)\b' <<< "${a_showconfig[*]}"; then
         a_output2+=(" - setting kernel module: \"$l_mod_name\" to
\"$(readlink -f /bin/false)\"")
         printf '%s\n' "install $l_mod_name $(readlink -f /bin/false)" >>
/etc/modprobe.d/"$l_mod_name".conf
      fi
      if ! grep -Pq -- '\bblacklist\h+'"${l_mod_name//-/_}"'\b' <<<
"${a_showconfig[*]}"; then
         a_output2+=(" - denylisting kernel module: \"$l_mod_name\"")
         printf '%s\n' "blacklist $l_mod_name" >>
/etc/modprobe.d/"$l_mod_name".conf
      fi
   }
   for l_mod_base_directory in $l_mod_path; do # Check if the module exists
on the system
      if [ -d "$l_mod_base_directory/${l_mod_name/-/\/}" ] && [ -n "$(ls -A
"$l_mod_base_directory/${l_mod_name/-/\/}")" ]; then
         a_output3+=("  - \"$l_mod_base_directory\"")
         [[ "$l_mod_name" =~ overlay ]] && l_mod_name="${l_mod_name::-2}"
         [ "$l_dl" != "y" ] && f_module_fix
      else
         echo -e " - kernel module: \"$l_mod_name\" doesn't exist in
\"$l_mod_base_directory\""
      fi
   done
   [ "${#a_output3[@]}" -gt 0 ] && printf '%s\n' "" " -- INFO --" " - module:
\"$l_mod_name\" exists in:" "${a_output3[@]}"
   [ "${#a_output2[@]}" -gt 0 ] && printf '%s\n' "" "${a_output2[@]}" ||
printf '%s\n' "" " - No changes needed"
   printf '%s\n' "" " - remediation of kernel module: \"$l_mod_name\"
complete" ""
}
```


Internal Only - General

References:

1. GRID: NIX-00000010
2. https://cve.mitre.org/cgi-bin/cvekey.cgi?keyword=filesystem

CIS Controls:

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

1.1.2 Configure Filesystem Partitions

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.

Using Btrfs subvolumes instead of traditional disk partitions changes how disk space is allocated and enforced, quotas may need to be configured to enforce storage limits and prevent resource exhaustion.

Note:

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

- Mount the new partition to a temporary mountpoint e.g. `mount /dev/sda2 /mnt`
- Copy data from the original partition to the new partition. e.g. `cp -a /var/tmp/* /mnt`
- Verify that all data is present on the new partition. e.g. `ls -la /mnt`

Internal Only - General

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

Internal Only - General

1.1.2.1 Configure /tmp

The `/tmp` directory is a world-writable directory used to store data used by the system and user applications for a short period of time. This data should have no expectation of surviving a reboot, as this directory is intended to be emptied after each reboot.

Internal Only - General

1.1.2.1.1 Ensure /tmp is tmpfs or a separate partition (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.

Note: If an entry for `/tmp` exists in `/etc/fstab` it will take precedence over entries in a systemd unit file. There is an exception to this when a system is diskless and connected to iSCSI, entries in `/etc/fstab` may not take precedence over a systemd unit file.

Rationale:

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

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

Impact:

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

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

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

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

Internal Only - General

Audit:

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

```
# findmnt -kn /tmp
```

Example output:

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

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

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

Example output:

```
generated
```

Verify output is not `masked` or `disabled`.

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

Remediation:

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

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

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

Example of using `tmpfs` with specific mount options:

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

Note: the `size=2G` is an example of setting a specific size for `tmpfs`.

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

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

Note: Using Btrfs subvolumes instead of traditional disk partitions changes how disk space is allocated and enforced, quotas may need to be configured to enforce storage limits and prevent resource exhaustion.

References:

1. GRID: NIX-00000043
2. CCI-000366
3. NIST SP 800-53 :: CM-6 b

Internal Only - General

4. NIST SP 800-53A :: CM-6.1 (iv)
5. STIG ID: RHEL-08-010543 | RULE ID: SV-230295r1017106 | CAT II
6. STIG ID: RHEL-09-231015 | RULE ID: SV-257844r1044918 | CAT II
7. STIG ID: ALMA-09-025540 | RULE ID: SV-269305r1050187 | CAT II
8. https://www.freedesktop.org/wiki/Software/systemd/APIFileSystems/
9. https://www.freedesktop.org/software/systemd/man/systemd-fstab-generator.html
10. https://www.kernel.org/doc/Documentation/filesystems/tmpfs.txt

Additional Information:

`/tmp` can be configured to use `tmpfs`.

`tmpfs` puts everything into the kernel internal caches and grows and shrinks to accommodate the files it contains and is able to swap unneeded pages out to swap space. It has maximum size limits which can be adjusted on the fly via `mount -o remount`.

Since `tmpfs` lives completely in the page cache and on swap, all `tmpfs` pages will be shown as "Shmem" in `/proc/meminfo` and "Shared" in `free`. Notice that these counters also include shared memory. The most reliable way to get the count is using `df` and `du`.

`tmpfs` has three mount options for sizing:

- `size`: The limit of allocated bytes for this `tmpfs` instance. The default is half of your physical RAM without swap. If you oversize your `tmpfs` instances the machine will deadlock since the OOM handler will not be able to free that memory.
- `nr_blocks`: The same as size, but in blocks of PAGE_SIZE.
- `nr_inodes`: The maximum number of inodes for this instance. The default is half of the number of your physical RAM pages, or (on a machine with highmem) the number of lowmem RAM pages, whichever is the lower.

These parameters accept a suffix k, m or g and can be changed on remount. The size parameter also accepts a suffix % to limit this `tmpfs` instance to that percentage of your physical RAM. The default, when neither `size` nor `nr_blocks` is specified, is `size=50%`.

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.	●	●	●

Internal Only - General

Controls Version	Control	IG 1	IG 2	IG 3
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.	●	●	●

Internal Only - General

1.1.2.1.2 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 create a block or character special devices in `/tmp`.

Audit:

- IF - a separate partition exists for `/tmp`, verify that the `nodev` option is set.

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

Example:

```
# findmnt -kn /tmp | grep -v nodev

Nothing should be returned
```

Remediation:

- IF - a separate partition exists for `/tmp`.

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

Example:

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

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

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

References:

1. GRID: NIX-00000044
2. NIST SP 800-53 :: CM-7 (2)
3. STIG ID: RHEL-08-040123 | RULE ID: SV-230511r958804 | CAT II
4. STIG ID: RHEL-09-231125 | RULE ID: SV-257866r958804 | CAT II
5. STIG ID: ALMA-09-026860 | RULE ID: SV-269317r1050199 | CAT II
6. CCI-001764

Internal Only - General

Additional Information:

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.	●	●	●

Internal Only - General

1.1.2.1.3 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:

- IF - a separate partition exists for `/tmp`, verify that the `nosuid` option is set.

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

Example:

```
# findmnt -kn /tmp | grep -v nosuid

Nothing should be returned
```

Remediation:

- IF - a separate partition exists for `/tmp`.

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

Example:

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

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

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

References:

1. GRID: NIX-00000045
2. NIST SP 800-53 :: CM-7 (2)
3. STIG ID: RHEL-08-040124 | RULE ID: SV-230512r958804 | CAT II
4. STIG ID: RHEL-09-231135 | RULE ID: SV-257868r958804 | CAT II
5. STIG ID: ALMA-09-027080 | RULE ID: SV-269319r1050201 | CAT II
6. CCI-001764

Internal Only - General

Additional Information:

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.	●	●	●

Internal Only - General

1.1.2.1.4 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`.

Impact:

Setting the `noexec` option on `/tmp` may prevent installation and/or updating of some 3rd party software.

Audit:

- IF - a separate partition exists for `/tmp`, verify that the `noexec` option is set.

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

Example:

```
# findmnt -kn /tmp | grep -v noexec

Nothing should be returned
```

Remediation:

- IF - a separate partition exists for `/tmp`.

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

Example:

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

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

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

References:

1. GRID: NIX-00000046

Internal Only - General

2. NIST SP 800-53 :: CM-7 (2)
3. STIG ID: RHEL-08-040125 | RULE ID: SV-230513r958804 | CAT II
4. STIG ID: RHEL-09-231130 | RULE ID: SV-257867r958804 | CAT II
5. STIG ID: ALMA-09-026970 | RULE ID: SV-269318r1050200 | CAT II
6. CCI-001764

Additional Information:

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.	●	●	●

Internal Only - General

1.1.2.2 Configure /dev/shm

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

Internal Only - General

1.1.2.2.1 Ensure /dev/shm is tmpfs or a separate partition (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

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

Rationale:

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

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

Impact:

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

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

Audit:

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

```
# findmnt -kn /dev/shm
```

Example output:

```
/dev/shm   tmpfs   tmpfs   rw,nosuid,nodev,noexec,relatime,seclabel
```

Remediation:

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

Internal Only - General

Example:

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

Note: Using Btrfs subvolumes instead of traditional disk partitions changes how disk space is allocated and enforced, quotas may need to be configured to enforce storage limits and prevent resource exhaustion.

References:

1. GRID: NIX-00000047
2. NIST SP 800-53 Rev. 5: CM-7
3. https://www.freedesktop.org/wiki/Software/systemd/APIFileSystems/
4. https://www.freedesktop.org/software/systemd/man/systemd-fstab-generator.html

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.	●	●	●

Internal Only - General

1.1.2.2.2 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:

- IF - a separate partition exists for `/dev/shm`, verify that the `nodev` option is set.

```
# findmnt -kn /dev/shm | grep -v 'nodev'

Nothing should be returned
```

Remediation:

- IF - a separate partition exists for `/dev/shm`.

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.

Example:

```
tmpfs /dev/shm    tmpfs     defaults,rw,nosuid,nodev,noexec,relatime  0 0
```

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

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

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

References:

1. GRID: NIX-00000048
2. NIST SP 800-53 :: CM-7 (2)
3. STIG ID: RHEL-08-040120 | RULE ID: SV-230508r958804 | CAT II
4. STIG ID: RHEL-09-231110 | RULE ID: SV-257863r958804 | CAT II
5. STIG ID: ALMA-09-026530 | RULE ID: SV-269314r1050196 | CAT II
6. CCI-001764

Internal Only - General

Additional Information:

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

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 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.	●	●	●

Internal Only - General

1.1.2.2.3 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:

- IF - a separate partition exists for `/dev/shm`, verify that the `nosuid` option is set.

```
# findmnt -kn /dev/shm | grep -v 'nosuid'

Nothing should be returned
```

Remediation:

- IF - a separate partition exists for `/dev/shm`.

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.

Example:

```
tmpfs /dev/shm    tmpfs     defaults,rw,nosuid,nodev,noexec,relatime  0 0
```

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

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

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

References:

1. GRID: NIX-00000049
2. NIST SP 800-53 :: CM-7 (2)
3. STIG ID: RHEL-08-040121 | RULE ID: SV-230509r958804 | CAT II
4. STIG ID: RHEL-09-231120 | RULE ID: SV-257865r1044946 | CAT II
5. STIG ID: ALMA-09-026750 | RULE ID: SV-269316r1050198 | CAT II
6. CCI-001764

Internal Only - General

Additional Information:

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

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 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.	●	●	●

Internal Only - General

1.1.2.2.4 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:

- IF - a separate partition exists for `/dev/shm`, verify that the `noexec` option is set.

```
# findmnt -kn /dev/shm | grep -v 'noexec'

Nothing should be returned
```

Remediation:

- IF - a separate partition exists for `/dev/shm`.

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

Example:

```
tmpfs /dev/shm    tmpfs     defaults,rw,nosuid,nodev,noexec,relatime  0 0
```

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

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

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

References:

1. GRID: NIX-00000050
2. NIST SP 800-53 :: CM-7 (2)
3. STIG ID: RHEL-08-040122 | RULE ID: SV-230510r958804 | CAT II
4. STIG ID: RHEL-09-231115 | RULE ID: SV-257864r958804 | CAT II

Internal Only - General

5. STIG ID: ALMA-09-026640 | RULE ID: SV-269315r1050197 | CAT II
6. CCI-001764

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.	●	●	●

Internal Only - General

1.1.2.3 Configure /home

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

Finding user home directories can be done by looking in `/etc/passwd`, looking over the mounted file systems with `mount` or querying the relevant database with `getent`.

The following script can be run to find user's home directories:

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

{
 l_valid_shells="^($(awk -F\/ '$NF != "nologin" {print}' /etc/shells | sed -
rn '/^\//{s,/,\\\\/,g;p}' | paste -s -d '|' - ))$"
 awk -v pat="$l_valid_shells" -F:
'($1!~/^(root|halt|sync|shutdown|nfsnobody)$/ && ($3>='"$(awk
'/^\s*UID_MIN/{print $2}' /etc/login.defs)"' || $3 != 65534) && $(NF) ~ pat)
{print $1 " - " $6}' /etc/passwd
}
```


Internal Only - General

1.1.2.3.1 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.

Using Btrfs subvolumes instead of traditional disk partitions changes how disk space is allocated and enforced, quotas may need to be configured to enforce storage limits and prevent resource exhaustion.

Rationale:

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

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

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

Impact:

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

Audit:

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

Internal Only - General

```
# findmnt -kn /home

/home   /dev/sdb   ext4   rw,nosuid,nodev,relatime,seclabel
```

Remediation:

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

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

Note: Using Btrfs subvolumes instead of traditional disk partitions changes how disk space is allocated and enforced, quotas may need to be configured to enforce storage limits and prevent resource exhaustion.

References:

1. GRID: NIX-00000051
2. NIST SP 800-53 Rev. 5: CM-7
3. AJ Lewis, "LVM HOWTO", http://tldp.org/HOWTO/LVM-HOWTO/

Additional Information:

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

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 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.	●	●	●

Internal Only - General

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

Audit:

- IF - a separate partition exists for `/home`, verify that the `nodev` option is set.

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

Example:

```
# findmnt -kn /home | grep -v nodev
```

Nothing should be returned

Remediation:

- IF - a separate partition exists for `/home`.

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

Example:

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

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

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

References:

1. GRID: NIX-00000052
2. NIST SP 800-53 Rev. 5: AC-3, MP-2
3. See the fstab(5) manual page for more information.

Internal Only - General

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.	●	●	●

Internal Only - General

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

Audit:

- IF - a separate partition exists for `/home`, verify that the `nosuid` option is set.

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

Example:

```
# findmnt -kn /home | grep -v nosuid
```

Nothing should be returned

Remediation:

- IF - a separate partition exists for `/home`.

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

Example:

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

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

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

References:

1. GRID: NIX-00000053
2. NIST SP 800-53 Rev. 5: AC-3, MP-2
3. See the fstab(5) manual page for more information.

Internal Only - General

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.	●	●	●

Internal Only - General

1.1.2.4 Configure /var

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

Internal Only - General

1.1.2.4.1 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:

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

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

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

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

Impact:

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

Audit:

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

Example:

Internal Only - General

```
# findmnt -kn /var

/var  /dev/sdb  ext4   rw,nosuid,nodev,relatime,seclabel
```

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.

Note: Using Btrfs subvolumes instead of traditional disk partitions changes how disk space is allocated and enforced, quotas may need to be configured to enforce storage limits and prevent resource exhaustion.

References:

1. GRID: NIX-00000054
2. NIST SP 800-53 Rev. 5: CM-7
3. RHEL 8 STIG Vul ID: V-244529
4. RHEL 8 STIG Rule ID: SV-244529r902737
5. 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 multi-user mode.

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.	●	●	●

Internal Only - General

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

Audit:

- IF - a separate partition exists for `/var`, verify that the `nodev` option is set.

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

Example:

```
# findmnt -kn /var | grep -v nodev
```

Nothing should be returned

Remediation:

- IF - a separate partition exists for `/var`.

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

Example:

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

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

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

References:

1. GRID: NIX-00000055
2. NIST SP 800-53 Rev. 5: AC-3, MP-2
3. See the fstab(5) manual page for more information.

Internal Only - General

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.	●	●	●

Internal Only - General

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

Audit:

- IF - a separate partition exists for `/var`, verify that the `nosuid` option is set.

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

Example:

```
# findmnt -kn /var | grep -v nosuid
```

Nothing should be returned

Remediation:

- IF - a separate partition exists for `/var`.

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

Example:

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

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

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

References:

1. GRID: NIX-00000056
2. NIST SP 800-53 Rev. 5: AC-3, MP-2
3. See the fstab(5) manual page for more information.

Internal Only - General

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.	●	●	●

Internal Only - General

1.1.2.5 Configure /var/tmp

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

Internal Only - General

1.1.2.5.1 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. Temporary files residing in `/var/tmp` are to be preserved between reboots.

Rationale:

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

Configuring `/var/tmp` as its own file system allows an administrator to set additional mount options such as `noexec/nosuid/nodev`. These options limit an attacker's ability to create exploits on the system.

Impact:

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

Audit:

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

Example:

```
# findmnt -kn /var/tmp
```

Example Output

```
/var/tmp   /dev/sdb ext4   rw,nosuid,nodev,noexec,relatime,seclabel
```

Remediation:

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

Internal Only - General

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

Note: Using Btrfs subvolumes instead of traditional disk partitions changes how disk space is allocated and enforced, quotas may need to be configured to enforce storage limits and prevent resource exhaustion.

References:

1. GRID: NIX-00000057
2. AJ Lewis, "LVM HOWTO", http://tldp.org/HOWTO/LVM-HOWTO/
3. NIST SP 800-53 Rev. 5: CM-7
4. NIST SP 800-53A :: CM-6.1 (iv)
5. RHEL 8 STIG Vul ID: V-244529
6. RHEL 8 STIG Rule ID: SV-244529r902737

Additional Information:

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

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

1.1.2.5.2 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 create a block or character special devices in `/var/tmp`.

Audit:

- IF - a separate partition exists for `/var/tmp`, verify that the `nodev` option is set.

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

Example:

```
# findmnt -kn /var/tmp | grep -v nodev

Nothing should be returned
```

Remediation:

- IF - a separate partition exists for `/var/tmp`.

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

Example:

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

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

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

References:

1. GRID: NIX-00000058
2. NIST SP 800-53 Rev. 5: AC-3, MP-2
3. NIST SP 800-53 Revision 5 :: CM-7 (2)
4. RHEL 8 STIG Vul ID: V-230520

Internal Only - General

5. RHEL 8 STIG Rule ID: SV-230520r854061
6. 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.	●	●	●

Internal Only - General

1.1.2.5.3 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:

- IF - a separate partition exists for `/var/tmp`, verify that the `nosuid` option is set.

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

Example:

```
# findmnt -kn /var/tmp | grep -v nosuid

Nothing should be returned
```

Remediation:

- IF - a separate partition exists for `/var/tmp`.

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

Example:

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

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

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

References:

1. GRID: NIX-00000059
2. NIST SP 800-53 Rev. 5: AC-3, MP-2
3. NIST SP 800-53 Revision 5 :: CM-7 (2)
4. RHEL 8 STIG Vul ID: V-230521

Internal Only - General

5. RHEL 8 STIG STIG ID: RHEL-08-040133
6. 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.	●	●	●

Internal Only - General

1.1.2.5.4 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:

- IF - a separate partition exists for `/var/tmp`, verify that the `noexec` option is set.

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

Example:

```
# findmnt -kn /var/tmp | grep -v noexec

Nothing should be returned
```

Remediation:

- IF - a separate partition exists for `/var/tmp`.

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

Example:

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

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

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

References:

1. GRID: NIX-00000060
2. NIST SP 800-53 Rev. 5: AC-3, MP-2
3. NIST SP 800-53 Revision 5 :: CM-7 (2)

Internal Only - General

4. RHEL 8 STIG Vul ID: V-230522
5. RHEL 8 STIG Rule ID: SV-230522r854063
6. 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.	●	●	●

Internal Only - General

1.1.2.6 Configure /var/log

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

Internal Only - General

1.1.2.6.1 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:

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

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

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

Impact:

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

Audit:

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

```
# findmnt -kn /var/log

/var/log /dev/sdb ext4   rw,nosuid,nodev,noexec,relatime,seclabel
```

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.

Internal Only - General

Note: Using Btrfs subvolumes instead of traditional disk partitions changes how disk space is allocated and enforced, quotas may need to be configured to enforce storage limits and prevent resource exhaustion.

References:

1. GRID: NIX-00000061
2. NIST SP 800-53 Rev. 5:: CM-6 b CM-7
3. NIST SP 800-53A :: CM-6.1 (iv)
4. RHEL 8 STIG Vul ID: V-230293
5. RHEL 8 STIG Rule ID: SV-230293r902720
6. 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 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.		●	●

Internal Only - General

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

Audit:

- IF - a separate partition exists for `/var/log`, verify that the `nodev` option is set.

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

Example:

```
# findmnt -kn /var/log | grep -v nodev

Nothing should be returned
```

Remediation:

- IF - a separate partition exists for `/var/log`.

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

Example:

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

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

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

References:

1. GRID: NIX-00000062
2. NIST SP 800-53 Rev. 5: AC-3, MP-2
3. NIST SP 800-53 Revision 5 :: CM-7 (2)
4. RHEL 8 STIG Vul ID: V-230514

Internal Only - General

5. RHEL 8 STIG Rule ID: SV-230514r854055
6. 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.	●	●	●

Internal Only - General

1.1.2.6.3 Ensure nosuid option set on /var/log 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/log` filesystem is only intended for log files, set this option to ensure that users cannot create `setuid` files in `/var/log`.

Audit:

- IF - a separate partition exists for `/var/log`, verify that the `nosuid` option is set.

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

Example:

```
# findmnt -kn /var/log | grep -v nosuid

Nothing should be returned
```

Remediation:

- IF - a separate partition exists for `/var/log`.

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

Example:

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

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

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

References:

1. GRID: NIX-00000063
2. NIST SP 800-53 Rev. 5: AC-3, MP-2
3. NIST SP 800-53 Revision 5 :: CM-7 (2)
4. RHEL 8 STIG Vul ID: V-230515

Internal Only - General

5. RHEL 8 STIG Rule ID: SV-230515r854056
6. 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.	●	●	●

Internal Only - General

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

Audit:

- IF - a separate partition exists for `/var/log`, verify that the `noexec` option is set.

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

Example:

```
# findmnt -kn /var/log | grep -v noexec

Nothing should be returned
```

Remediation:

- IF - a separate partition exists for `/var/log`.

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

Example:

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

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

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

References:

1. GRID: NIX-00000064
2. NIST SP 800-53 Rev. 5: AC-3, MP-2
3. NIST SP 800-53 Revision 5 :: CM-7 (2)

Internal Only - General

4. RHEL 8 STIG Vul ID: V-230516
5. RHEL 8 STIG Rule ID: SV-230516r854057
6. 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.	●	●	●

Internal Only - General

1.1.2.7 Configure /var/log/audit

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

Internal Only - General

1.1.2.7.1 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:

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

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

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

Impact:

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

Audit:

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

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

/var/log/audit /dev/sdb ext4   rw,nosuid,nodev,noexec,relatime,seclabel
```

Remediation:

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

Internal Only - General

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

Note: Using Btrfs subvolumes instead of traditional disk partitions changes how disk space is allocated and enforced, quotas may need to be configured to enforce storage limits and prevent resource exhaustion.

References:

1. GRID: NIX-00000065
2. NIST SP 800-53 Rev. 5: CM-7 CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. RHEL 8 STIG Vul ID: V-230294
5. RHEL 8 STIG Rule ID: SV-230294r627750
6. 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 multi-user mode.

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.		●	●

Internal Only - General

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

Audit:

- IF - a separate partition exists for `/var/log/audit`, verify that the `nodev` option is set.

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

Example:

```
# findmnt -kn /var/log/audit | grep -v nodev

Nothing should be returned
```

Remediation:

- IF - a separate partition exists for `/var/log/audit`.

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

Example:

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

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

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

References:

1. GRID: NIX-00000066
2. NIST SP 800-53 Revision 5 :: CM-7 (2)
3. STIG ID: RHEL-08-040129 | RULE ID: SV-230517r958804 | CAT II

Internal Only - General

4. STIG ID: RHEL-09-231175 | RULE ID: SV-257876r958804 | CAT II
5. STIG ID: ALMA-09-027190 | RULE ID: SV-269320r1050202 | CAT II
6. 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.	●	●	●

Internal Only - General

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

Audit:

- IF - a separate partition exists for `/var/log/audit`, verify that the `nosuid` option is set.

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

Example:

```
# findmnt -kn /var/log/audit | grep -v nosuid

Nothing should be returned
```

Remediation:

- IF - a separate partition exists for `/var/log/audit`.

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

Example:

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

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

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

References:

1. GRID: NIX-00000067
2. NIST SP 800-53 Rev. 5: AC-3, MP-2
3. NIST SP 800-53 Revision 5 :: CM-7 (2)

Internal Only - General

4. STIG ID: RHEL-08-040130 | RULE ID: SV-230518r958804 | CAT II
5. STIG ID: RHEL-09-231170 | RULE ID: SV-257875r958804 | CAT II
6. 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.	●	●	●

Internal Only - General

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

Audit:

- IF - a separate partition exists for `/var/log/audit`, verify that the `noexec` option is set.

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

Example:

```
# findmnt -kn /var/log/audit | grep -v noexec

Nothing should be returned
```

Remediation:

- IF - a separate partition exists for `/var/log/audit`.

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

Example:

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

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

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

References:

1. GRID: NIX-00000068
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

Internal Only - General

3. NIST SP 800-53 Revision 5 :: CM-7 (2)
4. STIG ID: RHEL-08-040131 | RULE ID: SV-230519r958804 | CAT II
5. STIG ID: RHEL-09-231165 | RULE ID: SV-257874r958804 | CAT II
6. 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.	●	●	●

1.2 Package Management

Patch management procedures may vary widely between enterprises. Large enterprises may choose to install a local updates server that can be used in place of their distributions servers, whereas a single deployment of a system may prefer to get updates directly. Updates can be performed automatically or manually, depending on the site's policy for patch management. Organizations may prefer to test patches against their environment on a non-production system before rolling out to production.

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

Software updates often offer new and improved features and speed enhancements.

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

Internal Only - General

1.2.1 Configure Package Repositories

Patch management procedures may vary widely between enterprises. Large enterprises may choose to install a local updates server that can be used in place of their distributions servers, whereas a single deployment of a system may prefer to get updates directly. Updates can be performed automatically or manually, depending on the site's policy for patch management. Organizations may prefer to test patches against their environment on a non-production system before rolling out to production.

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

Note: Creation of an appropriate patch management policy is left to the organization.

Internal Only - General

1.2.1.1 Ensure the source.list and .source files use the Signed-By option (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`sources.list` - List of configured APT data sources

The source list `/etc/apt/sources.list` and the files contained in `/etc/apt/sources.list.d/` are designed to support any number of active sources and a variety of source media.

Signed-By (`signed-by`) is an option to require a repository to pass apt-secure(8) verification with a certain set of keys rather than all trusted keys apt has configured. It is specified as a list of absolute paths to keyring files (have to be accessible and readable for the _apt system user, so ensure everyone has read-permissions on the file) and fingerprints of keys to select from these keyrings. The recommended locations for keyrings are `/usr/share/keyrings` for keyrings managed by packages, and `/etc/apt/keyrings` for keyrings managed by the system operator. If no keyring files are specified, the default is the `trusted.gpg` keyring and all keyrings in the `trusted.gpg.d/` directory. If no fingerprint is specified all keys in the keyrings are selected. A fingerprint will accept also all signatures by a subkey of this key, if this isn't desired an exclamation mark (`!`) can be appended to the fingerprint to disable this behaviour. The option defaults to the value of the option with the same name if set in the previously acquired Release file of this repository (only fingerprints can be specified there through). Otherwise, all keys in the trusted keyrings are considered valid signers for this repository. The option may also be set directly to an embedded GPG public key block.

Rationale:

Verifying that updates originate from a trusted and authenticated source helps mitigate the risk of spoofing attacks, which could otherwise lead to the inadvertent installation of malicious or compromised software.

Impact:

`apt-key list` is deprecated. Manage keyring files in `trusted.gpg.d` instead (see apt-key(8)). With the deprecation of `apt-key` it is recommended to use the `Signed-By` option in `sources.list` to require a repository to pass apt-secure(8) verification with a certain set of keys rather than all trusted keys apt has configured.

Internal Only - General

It is important to list sources in order of preference, with the most preferred source listed first. Typically, this will result in sorting by speed from fastest to slowest (CD-ROM followed by hosts on a local network, followed by distant Internet hosts, for example).

Audit:

Step 1 Legacy `.list` format (per-line check)

Run the following command to find uncommented `deb`/`deb-src` entries that lack a `Signed-By` option:

```
# grep -PRn -- '^\h*(deb(-src)?)\h+(\[(?![^\]]*\bSigned-By\b)[^\]]*\]|[^\[])'
\
    /etc/apt/sources.list \
    /etc/apt/sources.list.d/*.list 2>/dev/null
```

Nothing should be returned.

> Note: This command checks each `deb`/`deb-src` line individually. It reports lines that either have no options bracket at all, or have an options bracket that does not include `Signed-By`. Commented lines (`#`) are excluded.

Step 2 deb822 `.sources` format (per-stanza check)

Run the following script to find stanzas that have a `URIs:` field but no `Signed-By:` field:

```
#!/usr/bin/env bash
{
  shopt -s nullglob
  while IFS= read -r -d $'\0' l_file; do
    awk '
      BEGIN { uris=0; signed=0; lineno=0 }
      /^[[:blank:]]*$/ {
        if (uris && !signed)
          print FILENAME ": stanza ending near line " lineno " is missing
Signed-By"
        uris=0; signed=0
      }
      /^[[:blank:]]*URIs:[[:blank:]]*\S/ { uris=1 }
      /^[[:blank:]]*Signed-By:[[:blank:]]*\S/ { signed=1 }
      { lineno=NR }
      END { if (uris && !signed)
          print FILENAME ": final stanza missing Signed-By" }
    ' "$l_file"
  done < <(find /etc/apt/sources.list.d/ -maxdepth 1 -type f -name
'*.sources' -print0)
}
```

Nothing should be returned.

Remediation:

Update the file `sources.list` and any `.sources` files found in the `/etc/apt/sources.list.d/` directory in accordance with site policy.

Internal Only - General

Note:

- The suggested filename for new systems is `/etc/apt/sources.list.d/vendor.sources`, where vendor is the result of `dpkg-vendor --query Vendor | tr A-Z a-z`, in deb822-style format. For example, Ubuntu uses `/etc/apt/sources.list.d/ubuntu.sources`.
- It is important to list sources in order of preference, with the most preferred source listed first. Typically, this will result in sorting by speed from fastest to slowest (CD-ROM followed by hosts on a local network, followed by distant Internet hosts, for example).

References:

1. GRID: NIX-00000080
2. NIST SP 800-53 Rev. 5: SI-2
3. https://manpages.debian.org/trixie/apt/sources.list.5.en.html
4. https://manpages.debian.org/trixie/gpg/gpg.1.en.html

CIS Controls:

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

Internal Only - General

1.2.1.2 Ensure weak dependencies are configured (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

The Debian packaging specification allows for 3 levels of dependencies to be declared:

hard dependencies (`Requires`) for packages which must be installed for a minimal installation of the application to function weak dependencies (`Recommends`) for packages which provide additional features, but which are not required for a minimal installation to function hints (`Suggests`) for packages which offer add-ons which might be useful.

Rationale:

Unless a system specifically requires the additional capabilities provides by the weak dependencies, it is recommended that the packages are not installed to reduce the potential attack surface.

Audit:

Run the following command and verify that the APT options for weak dependencies are disabled:

```
# /usr/bin/apt-config dump | /usr/bin/grep "APT::Install-"
```

Verify the output matches the following

```
APT::Install-Recommends "0";
APT::Install-Suggests "0";
```

Note: If no output is returned, this should be considered a failure as the default behavior is to install packages marked as recommended.

Remediation:

Create a file under `/etc/apt/apt.conf.d/` to disable the installation of weak dependencies:

```
# printf '%s\n%s\n' "" 'APT::Install-Recommends "0";' 'APT::Install-Suggests
"0";' > /etc/apt/apt.conf.d/60-no-weak-dependencies
```

Note: All configuration files under `/etc/apt/apt.conf.d/` are sorted in lexicographic order and where conflicting settings exist the last one will take precedence.

Internal Only - General

Default Value:

```
APT::Install-Recommends "1";
APT::Install-Suggests "0";
```

References:

1. GRID: NIX-00000082
2. NIST SP 800-53::SI-2

CIS Controls:

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

Internal Only - General

1.2.1.3 Ensure access to gpg key files are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

GPG key files are cryptographic keys used by APT to verify the authenticity and integrity of software packages. Each key ensures that packages come from a trusted source and have not been tampered with. These keys are typically stored in keyring directories `/etc/apt/trusted.gpg.d` or `/usr/share/keyrings` and referenced in repository configuration files using the `signed-by` option for secure package verification.

Rationale:

GPG key files are the trust anchors APT relies on to verify that downloaded packages and repository metadata have not been tampered with. If a non-root user can write to these key files, they could replace or add a key that vouches for a malicious or spoofed repository, causing APT to accept and install compromised packages as if they were legitimate. Restricting access to root ensures only trusted administrative processes can change the keys accepted for package verification.

Impact:

The `apt-key` command is deprecated, and you should manage keyring files individually in the `trusted.gpg.d` directory instead. This change is due to security concerns with the way `apt-key` works, as it adds keys to the `/etc/apt/trusted.gpg` file, which can lead to potential security issues.

Audit:

1. Run the following script to verify all `.gpg` key files in `/usr/share/keyrings/` and `/etc/apt/trusted.gpg.d` are mode `0644` or more restrictive, are owned by the user `root`, and group owned by the group `root`:

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

{
   while IFS= read -r -d $'\0' l_file; do
      stat -Lc 'File: %n Mode: (%#a) User: (%U) Group: (%G)' "$l_file"
   done < <(find -L /usr/share/keyrings/ /etc/apt/trusted.gpg.d/ \
   -mount -xdev -type f \( ! -user root -o ! -group root -o -perm /133 \) \
   -name '*gpg' -print0)
}
```

Nothing should be returned.

Internal Only - General

2. Run the following script to verify that `Signed-By` key files referenced in `.list` and `.sources` source files, that reside outside the standard keyring locations already covered by Step 1, are mode `0644` or more restrictive, are owned by the user `root`, and group owned by the group `root`:

```
#!/usr/bin/env bash
{
   while IFS= read -r l_keypath; do
      l_keyfile="$(readlink -e "$l_keypath" 2>/dev/null)"
      [ -z "$l_keyfile" ] && continue
      case "$l_keyfile" in
        /usr/share/keyrings/*|/etc/apt/trusted.gpg.d/*) continue ;;
      esac
      stat -Lc 'File: %n Mode: (%#a) User: (%U) Group: (%G)' "$l_keyfile" | \
      grep -Pv -- 'Mode: \(0?6[04][04]\) User: \(root\) Group: \(root\)'
   done < <(grep -PRho -- '\bSigned-By\b\s*[=:]\s*\K\S+' \
      /etc/apt/sources.list \
      /etc/apt/sources.list.d/*.list \
      /etc/apt/sources.list.d/*.sources 2>/dev/null | sort -u)
}
```

Nothing should be returned.

> **Note:** This step audits key files referenced by `Signed-By=` entries that reside outside the standard keyring locations already covered by Step 1. Common alternative locations include `/etc/keyrings/` and `/usr/local/share/keyrings/`.

Remediation:

Run the following commands to update the access to any file(s) returned in the audit steps above:

```
# chown root:root "/absolute/path/to/file"
# chmod  u-x,go-wx "/absolute/path/to/file"
```

Example

```
# chown root:root /usr/share/keyrings/debian-archive-trixie-security-
automatic.gpg
# chmod  u-x,go-wx /usr/share/keyrings/debian-archive-trixie-security-
automatic.gpg
```

Note: APT first checks keys specified with the `signed-by` option in a repository's `.sources` or `.list` files. If no key is specified there, it falls back to administrator-managed keyrings in `/etc/apt/trusted.gpg.d`, then system-wide keyrings in `/usr/share/keyrings`.

Internal Only - General

Default Value:

```
Access: (644/-rw-r--r--) Uid: ( 0/ root) Gid: ( 0/ root)
```

References:

1. GRID: NIX-00000075
2. https://ftp-master.debian.org/keys.html
3. https://wiki.debian.org/SecureApt#All_about_secure_apt
4. https://manpages.debian.org/stretch/apt/apt-secure.8.en.html

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.	●	●	●

Internal Only - General

1.2.1.4 Ensure access to /etc/apt/trusted.gpg.d directory is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `/etc/apt/trusted.gpg.d` directory stores individual GPG key files that APT uses to verify the authenticity of packages and repositories.

Rationale:

Impact:

A non-privileged user with write access to `/etc/apt/trusted.gpg.d` can compromise the chain of trust by adding new public keys to APT trusted keys configuration.

Audit:

Run the following command and verify `/etc/apt/trusted.gpg.d` `Uid` and `Gid` are both `0/root` and access is set to `0755` or more secure

```
# stat -Lc 'Access: (%a/%A) Uid: ( %u/ %U) Gid: ( %g/ %G)'
/etc/apt/trusted.gpg.d

Access: (755/drwxr-xr-x) Uid: ( 0/ root) Gid: ( 0/ root)
```

Remediation:

Run the following commands to set `/etc/apt/trusted.gpg.d` `Uid` and `Gid` to `0/root` and access to `0755` or more restrictive:

```
# chown root:root /etc/apt/trusted.gpg.d
# chmod u=rwx,g=rx,o=rx /etc/apt/trusted.gpg.d
```

Default Value:

Access: (755/drwxr-xr-x) Uid: (0/ root) Gid: (0/ root)

References:

1. GRID: NIX-00000071
2. NIST SP 800-53 Rev. 5: AC-3, MP-2
3. NIST SP 800-53A::SC-5 (2).1

Internal Only - General

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.	●	●	●

Internal Only - General

1.2.1.5 Ensure access to /etc/apt/auth.conf.d directory is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`/etc/apt/auth.conf.d` contains `.conf` files with credentials to access APT private repositories.

Rationale:

The `/etc/apt/auth.conf.d` directory contains APT `.conf` files which should have restricted access because they control how the package manager runs as root and may contain sensitive information. If writable or readable by non-root users, they could be altered to run malicious commands or expose credentials.

Audit:

Run the following command and verify `/etc/apt/auth.conf.d` `Uid` and `Gid` are both `0/root` and access is `0755` or more restrictive:

```
# stat -Lc 'Access: (%a/%A) Uid: ( %u/ %U) Gid: ( %g/ %G)'
/etc/apt/auth.conf.d

Access: (755/drwxr-xr-x) Uid: ( 0/ root) Gid: ( 0/ root)
```

Remediation:

Run the following commands to set permissions on `/etc/apt/auth.conf.d/` `Uid` and `Gid` are both `0/root` and access is `0755` or more restrictive:

```
# chown root:root /etc/apt/auth.conf.d
# chmod u=rwx,g=rx,o=rx /etc/apt/auth.conf.d
```

Default Value:

Access: (755/drwxr-xr-x) Uid: (0/ root) Gid: (0/ root)

References:

1. GRID: NIX-00000069
2. NIST SP 800-53 Rev. 5: AC-3, MP-2
3. NIST SP 800-53A::SC-5 (2).1
4. https://manpages.debian.org/trixie/apt/apt_auth.conf.5.en.html

Internal Only - General

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.	●	●	●

Internal Only - General

1.2.1.6 Ensure access to files in the /etc/apt/auth.conf.d/ directory is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Files in the `/etc/apt/auth.conf.d` directory contains `.conf` files with configured credentials for private APT repositories or proxy.

Rationale:

Access to the `.conf` files in the `/etc/apt/auth.conf.d` directory should be restricted to prevent secrets leakage of Login information needed to connect to a proxy or to download data from a repository.

Audit:

Run the following command and verify all `.conf` files in the `/etc/apt/auth.conf.d/` directory `Uid` and `Gid` are both `0/root` and mode `0640` or more restrictive:

Run the following command:

```
# stat -Lc 'Access: (%a/%A) Uid: ( %u/ %U) Gid: ( %g/ %G)'
/etc/apt/auth.conf.d/* 2>/dev/null
```

No output should be returned.

Remediation:

Run the following commands to set all `.conf` files in the `/etc/apt/auth.conf.d/` directory `Uid` and `Gid` both to`0/root` and mode `0640`:

```
# chown root:root /etc/apt/auth.conf.d/*
# chmod u-x,g-wx,o-rwx /etc/apt/auth.conf.d/*
```

Default Value:

The `/etc/apt/auth.conf.d/` directory is empty

References:

1. GRID: NIX-00000074
2. NIST SP 800-53 Rev. 5: AC-3, MP-2
3. NIST SP 800-53A::SC-5 (2).1
4. https://manpages.debian.org/testing/apt/apt_auth.conf.5.en.html

Internal Only - General

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.	●	●	●

Internal Only - General

1.2.1.7 Ensure access to /usr/share/keyrings directory is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `/usr/share/keyrings` directory contains trusted gpg keyring files used by APT to verify package signatures from repositories. It provides a centralized, read-only location for system-wide keys that ensures package integrity and authenticity.

Rationale:

A non-root user should not be able to add or remove keys from this directory. Misconfiguring could allow a non-root user to add repositories containing malicious packages.

Audit:

Run the following command and verify `/usr/share/keyrings` `Uid` and `Gid` are both `0/root` and access is `0755` or more restrictive:

```
# stat -Lc 'Access: (%a/%A) Uid: ( %u/ %U) Gid: ( %g/ %G)'
/usr/share/keyrings

Access: (755/drwxr-xr-x) Uid: ( 0/ root) Gid: ( 0/ root)
```

Remediation:

Run the following commands to set permissions on `/usr/share/keyrings` `Uid` and `Gid` both `0/root` and access is `0755` or more restrictive:

```
# chown root:root /usr/share/keyrings
# chmod u=rwx,g=rx,o=rx /usr/share/keyrings
```

Default Value:

Access: (755/drwxr-xr-x) Uid: (0/ root) Gid: (0/ root)

References:

1. GRID: NIX-00000072

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	14.6 Train Workforce Members on Recognizing and Reporting Security Incidents Train workforce members to be able to recognize a potential incident and be able to report such an incident.	●	●	●
v7	3.3 Protect Dedicated Assessment Accounts Use a dedicated account for authenticated vulnerability scans, which should not be used for any other administrative activities and should be tied to specific machines at specific IP addresses.		●	●

Internal Only - General

1.2.1.8 Ensure access to /etc/apt/sources.list.d directory is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `/etc/apt/sources.list.d` directory provides a way to add `sources.list` entries in separate files.

Rationale:

`/etc/apt/sources.list.d` directory contains files with repositories to be used by APT. A non-root user should not be able to add or remove files from this directory. Misconfiguring could allow a non-root user to add repositories containing malicious packages.

Audit:

Run the following command and verify `/etc/apt/sources.list.d` `Uid` and `Gid` are both `0/root` and access is set to `0755` or more restrictive:

```
# stat -Lc 'Access: (%a/%A) Uid: ( %u/ %U) Gid: ( %g/ %G)'
/etc/apt/sources.list.d

Access: (755/drwxr-xr-x) Uid: ( 0/ root) Gid: ( 0/ root)
```

Remediation:

Run the following command to set permissions to `/etc/apt/sources.list.d` `Uid` and `Gid` to `0/root` and access to `0755` or more restrictive:

```
# chown root:root /etc/apt/sources.list.d
# chmod u=rwx,g=rx,o=rx /etc/apt/sources.list.d
```

Default Value:

Access: (755/drwxr-xr-x) Uid: (0/ root) Gid: (0/ root)

References:

1. GRID: NIX-00000070
2. https://manpages.debian.org/stretch/apt/sources.list.5.en.html

Internal Only - General

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.	●	●	●

Internal Only - General

1.2.1.9 Ensure access to files in /etc/apt/sources.list.d are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Files in the `/etc/apt/sources.list.d` directory provides a way to add `sources.list` entries in separate files.

Rationale:

Files in the `/etc/apt/sources.list.d` directory contain information about repositories to be used by APT. A non-root user should not be able to add or remove files from this directory. Misconfiguring could allow a non-root user to add repositories containing malicious packages.

Audit:

Run the following command and verify all files in the `/etc/apt/sources.list.d` directory `Uid` and `Gid` are both `0/root` and access is set to `0644` or more restrictive:

```
# stat -Lc 'Access: (%a/%A) Uid: ( %u/ %U) Gid: ( %g/ %G)'
/etc/apt/sources.list.d/* 2>/dev/null
```

Example Output

```
Access: (644/-rw-r--r--) Uid: ( 0/ root) Gid: ( 0/ root)
```

Remediation:

Run the following command to set permissions to `/etc/apt/sources.list.d` directory `Uid` and `Gid` to `0/root` and access to `0644` or more restrictive:

```
# chown root:root /etc/apt/sources.list.d/*
# chmod u-x,go-wx /etc/apt/sources.list.d/*
```

Default Value:

No files under /etc/apt/sources.list.d

References:

1. GRID: NIX-00000073
2. https://manpages.debian.org/stretch/apt/sources.list.5.en.html

Internal Only - General

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.	●	●	●

Internal Only - General

1.2.1.10 Ensure repositories are set to https in /etc/apt/sources.list (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

APT supports both HTTP and HTTPS transport for repository access. While official repositories for the target distributions use HTTPS by default, custom or third-party repositories added by administrators may be configured with HTTP. Repositories served over plain HTTP are susceptible to network interception — an attacker positioned between the system and a mirror could suppress specific security updates, withholding patches to keep the system vulnerable to known exploits, or intercept repository metadata to profile installed software versions.

Rationale:

APT verifies package integrity via GPG signatures, but plain HTTP transport leaves repository metadata and package lists exposed to network interception. An attacker positioned between the system and a repository mirror could suppress specific security updates, preventing patches from being applied and keeping the system vulnerable to known exploits. Enforcing HTTPS transport ensures the integrity and confidentiality of the repository communication channel, complementing GPG signature verification.

Impact:

Changing a repository from HTTP to HTTPS requires that the repository server supports HTTPS. For local mirrors that serve only HTTP, switching the client to HTTPS without a corresponding server-side change will cause `apt update` to fail for that repository.

Audit:

Run the following commands to verify no repositories in `/etc/apt/sources.list` are configured to use plain HTTP:

1. Check legacy `.list` format:

```
# grep -Pri -- '^\s*deb\s+http://' /etc/apt/sources.list 2>/dev/null
```

Nothing should be returned.

Remediation:

Edit `/etc/apt/sources.list` and replace `http://` with `https://` for each repository URI.

Internal Only - General

Example — legacy `.list` format:

```
deb https://deb.debian.org/debian bookworm main
```

Default Value:

Official Debian 12 and Ubuntu 22.04+ repositories use HTTPS by default. Third-party or custom repositories may use HTTP.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.10 Encrypt Sensitive Data in Transit Encrypt sensitive data in transit. Example implementations can include: Transport Layer Security (TLS) and Open Secure Shell (OpenSSH).		●	●
v7	14.4 Encrypt All Sensitive Information in Transit Encrypt all sensitive information in transit.		●	●

Internal Only - General

1.2.1.11 Ensure repositories are set to https in /etc/apt/sources.list.d directory (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

APT supports both HTTP and HTTPS transport for repository access. While official repositories for the target distributions use HTTPS by default, custom or third-party repositories added by administrators may be configured with HTTP. Repositories served over plain HTTP are susceptible to network interception — an attacker positioned between the system and a mirror could suppress specific security updates, withholding patches to keep the system vulnerable to known exploits, or intercept repository metadata to profile installed software versions.

Rationale:

APT verifies package integrity via GPG signatures, but plain HTTP transport leaves repository metadata and package lists exposed to network interception. An attacker positioned between the system and a repository mirror could suppress specific security updates, preventing patches from being applied and keeping the system vulnerable to known exploits. Enforcing HTTPS transport ensures the integrity and confidentiality of the repository communication channel, complementing GPG signature verification.

Impact:

Changing a repository from HTTP to HTTPS requires that the repository server supports HTTPS. For local mirrors that serve only HTTP, switching the client to HTTPS without a corresponding server-side change will cause `apt update` to fail for that repository.

Audit:

Run the following commands to verify no repository files in `/etc/apt/sources.list.d/` are configured to use plain HTTP:

1. Check legacy `.list` format files:

```
# grep -Pri -- '^\s*deb\s+http://' /etc/apt/sources.list.d/*.list 2>/dev/null
```

Nothing should be returned.

2. Check deb822 `.sources` format files:

Internal Only - General

```
# grep -Pri -- '^\s*URIs:\s+http://' /etc/apt/sources.list.d/*.sources
2>/dev/null
```

Nothing should be returned.

Remediation:

Edit each file in `/etc/apt/sources.list.d/` that contains HTTP repository URIs and replace `http://` with `https://`.

Example — legacy `.list` format:

```
deb https://deb.debian.org/debian bookworm main
```

Example — deb822 `.sources` format:

```
URIs: https://deb.debian.org/debian
```

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.10 Encrypt Sensitive Data in Transit Encrypt sensitive data in transit. Example implementations can include: Transport Layer Security (TLS) and Open Secure Shell (OpenSSH).		●	●
v7	14.4 Encrypt All Sensitive Information in Transit Encrypt all sensitive information in transit.		●	●

Internal Only - General

1.2.1.12 Ensure AllowInsecureRepositories is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The APT configuration option `Acquire::AllowInsecureRepositories` controls whether APT is permitted to access repositories that do not have a valid cryptographic signature on their `InRelease` or `Release.gpg` metadata files. When set to `1`, APT bypasses repository authentication and will download package lists and packages from unsigned or improperly signed repositories without warning. This option should be explicitly set to `0` to enforce repository signature verification.

Rationale:

Allowing insecure repositories removes a critical layer of the APT trust chain. An attacker who can perform a man-in-the-middle attack, compromise a mirror, or redirect DNS could serve unsigned repository metadata and packages that APT would accept without verification. This could result in the installation of malicious or backdoored packages indistinguishable from legitimate ones. Explicitly setting `AllowInsecureRepositories "0"` ensures APT refuses to proceed when repository signatures are absent or invalid.

Impact:

Setting `AllowInsecureRepositories "0"` will cause `apt update` to fail for any repository that does not provide a valid signed `InRelease` or `Release.gpg` file. If a legitimate repository in use does not support signing, it must be removed or replaced with a signed alternative before applying this setting. Internal or private repositories must be configured to sign their metadata.

Audit:

Run the following command to verify `AllowInsecureRepositories` is set to `0`:

```
# apt-config dump | grep AllowInsecureRepositories
```

Verify output includes:

```
Acquire::AllowInsecureRepositories "0";
```

Remediation:

Create a configuration file to explicitly disable insecure repository access:

Internal Only - General

```
# printf '%s\n' "" 'Acquire::AllowInsecureRepositories "0";' >>
/etc/apt/apt.conf.d/99-no-insecure-repositories
```

Verify the setting has taken effect:

```
# apt-config dump | grep AllowInsecureRepositories
```

Note: All configuration files under `/etc/apt/apt.conf.d/` are processed in lexicographic order. Where conflicting settings exist, the last one takes precedence. The `99-` prefix ensures this setting overrides any earlier configuration.

Default Value:

Acquire::AllowInsecureRepositories "0";

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	7.3 Perform Automated Operating System Patch Management Perform operating system updates on enterprise assets through automated patch management on a monthly, or more frequent, basis.	●	●	●
v7	3.4 Deploy Automated Operating System Patch Management Tools 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.	●	●	●

Internal Only - General

1.2.1.13 Ensure AllowWeakRepositories is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The APT configuration option `Acquire::AllowWeakRepositories` controls whether APT accepts repositories signed with cryptographically weak algorithms. Weak signatures include those using MD5 or SHA-1 hashing algorithms, which are no longer considered cryptographically secure. When set to `1`, APT treats repositories signed with these deprecated algorithms as trusted. This option should be explicitly set to `0` to enforce strong signature requirements.

Rationale:

MD5 and SHA-1 are cryptographically broken hash algorithms for which practical collision and preimage attacks exist. A repository signed with a weak algorithm cannot be trusted to authenticate the integrity of package metadata. An attacker who can produce a collision could craft malicious repository metadata that passes signature verification, enabling installation of tampered packages. Explicitly setting `AllowWeakRepositories "0"` ensures APT rejects repositories using deprecated cryptographic algorithms.

Impact:

Setting `AllowWeakRepositories "0"` will cause `apt update` to fail for any repository signed exclusively with weak algorithms (MD5 or SHA-1). Repositories using these algorithms must be migrated to strong signing keys (RSA-4096 or Ed25519 with SHA-256 or SHA-512) before this setting can be applied without disruption.

Audit:

Run the following command to verify `AllowWeakRepositories` is set to `0`:

```
# apt-config dump | grep AllowWeakRepositories
```

Verify output includes:

```
Acquire::AllowWeakRepositories "0";
```

Remediation:

Create or append to a configuration file to explicitly disable weak repository signatures:

Internal Only - General

```
# printf '%s\n' "" 'Acquire::AllowWeakRepositories "0";' >>
/etc/apt/apt.conf.d/99-no-insecure-repositories
```

Verify the setting has taken effect:

```
# apt-config dump | grep AllowWeakRepositories
```

Note: All configuration files under `/etc/apt/apt.conf.d/` are processed in lexicographic order. Where conflicting settings exist, the last one takes precedence.

Default Value:

Acquire::AllowWeakRepositories "0";

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	7.3 Perform Automated Operating System Patch Management Perform operating system updates on enterprise assets through automated patch management on a monthly, or more frequent, basis.	●	●	●
v7	3.4 Deploy Automated Operating System Patch Management Tools 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.	●	●	●

Internal Only - General

1.2.1.14 Ensure AllowDowngradeToInsecureRepositories is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The APT configuration option `Acquire::AllowDowngradeToInsecureRepositories` controls whether APT will silently fall back to an insecure (unsigned or weakly signed) repository when a secure version of the same repository is unavailable. When set to `1`, APT permits this downgrade without user confirmation. This option should be explicitly set to `0` to prevent silent downgrade attacks.

Rationale:

Allowing silent downgrade to insecure repositories creates a downgrade attack vector. An attacker who can block access to the HTTPS or signed version of a repository (through network interference, DNS poisoning, or mirror compromise) could force APT to fall back to an insecure alternative under their control. This would allow serving malicious packages that APT accepts without warning. Explicitly setting `AllowDowngradeToInsecureRepositories "0"` ensures APT fails loudly rather than silently accepting an insecure fallback.

Impact:

Setting `AllowDowngradeToInsecureRepositories "0"` means that if a secure repository becomes temporarily unavailable and an insecure fallback is configured, `apt update` will fail rather than falling back silently. This is the safer failure mode — administrators will be alerted to the unavailability of the secure repository rather than unknowingly operating against an insecure one.

Audit:

Run the following command to verify `AllowDowngradeToInsecureRepositories` is set to `0`

```
# apt-config dump | grep AllowDowngradeToInsecureRepositories
```

Verify output includes:

```
Acquire::AllowDowngradeToInsecureRepositories "0";
```

Remediation:

Create or append to a configuration file to explicitly disable insecure repository downgrade:

Internal Only - General

```
# echo 'Acquire::AllowDowngradeToInsecureRepositories "0";' >>
/etc/apt/apt.conf.d/99-no-insecure-repositories
```

Verify the setting has taken effect:

```
# apt-config dump | grep AllowDowngradeToInsecureRepositories
```

Note: All configuration files under `/etc/apt/apt.conf.d/` are processed in lexicographic order. Where conflicting settings exist, the last one takes precedence.

Default Value:

Acquire::AllowDowngradeToInsecureRepositories "0";

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	7.3 Perform Automated Operating System Patch Management Perform operating system updates on enterprise assets through automated patch management on a monthly, or more frequent, basis.	●	●	●
v7	3.4 Deploy Automated Operating System Patch Management Tools 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.	●	●	●

Internal Only - General

1.2.1.15 Ensure Acquire::Check-Date is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The APT configuration option `Acquire::Check-Date` controls whether APT verifies that the system clock is plausible relative to the `Valid-Until` field in a repository's `Release` file. When enabled (the default), APT will reject `Release` files dated significantly in the future or past, preventing the use of stale repository metadata. This option should be explicitly set to `true` to ensure this validation is always active and cannot be silently disabled by other configuration.

Rationale:

Repository `Release` files include a `Valid-Until` field specifying a date after which the metadata should be considered stale. A man-in-the-middle attacker who intercepts APT traffic could serve an older `Release` file to make the system appear up-to-date while pinning it to outdated, vulnerable package versions. Without clock validity checking, APT may accept stale metadata and report no available updates, leaving known vulnerabilities unpatched. Explicitly enabling `Acquire::Check-Date` prevents this class of replay attack against the APT update process.

Impact:

If the system clock is significantly incorrect (e.g., due to misconfigured NTP, hardware clock drift, or clock manipulation), enabling `Check-Date` will cause `apt update` to fail with a date-related error. Ensure NTP synchronization is correctly configured (see the Time Synchronization section of the benchmark) before applying this setting. This is a fail-safe behavior that prevents operating against potentially stale repository metadata.

Audit:

Run the following command to verify `Acquire::Check-Date` is explicitly set to `true`:

```
# apt-config dump | grep -Psi -- 'Acquire::Check-Date'
```

Verify the output includes:

```
Acquire::Check-Date "true";
```

Note: If no output is returned, APT is using its compiled-in default (which is `true` on current Debian and Ubuntu releases), but explicit configuration is preferred to prevent override by other configuration files.

Internal Only - General

Remediation:

Create or append to a configuration file to explicitly enable date checking:

```
# printf '%s\n' "" 'Acquire::Check-Date "true";' >> /etc/apt/apt.conf.d/99-no-insecure-repositories
```

Verify the setting has taken effect:

```
# apt-config dump | grep -Psi -- 'Acquire::Check-Date'
```

Note: All configuration files under `/etc/apt/apt.conf.d/` are processed in lexicographic order. Where conflicting settings exist, the last one takes precedence.

Default Value:

Acquire::Check-Date "true";

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.4 Enforce Data Retention Retain data according to the enterprise's data management process. Data retention must include both minimum and maximum timelines.	●	●	●
v7	7.3 Limit Use of Scripting Languages in Web Browsers and Email Clients Ensure that only authorized scripting languages are able to run in all web browsers and email clients.		●	●

Internal Only - General

1.2.2 Configure Package Updates

Internal Only - General

1.2.2.1 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:

Verify there are no updates or patches to install:

```
# apt update
# apt -s upgrade
```

Check to make sure no system reboot is required:

```
# if [ -f /var/run/reboot-required ]; then
     cat /var/run/reboot-required
  fi
```

Remediation:

Run the following commands to update all packages following local site policy guidance on applying updates and patches:

Run the following command to refresh the local package index from configured repositories:

```
# apt update
```

Run one of the following commands to download and install all available updates and patches:

```
# apt upgrade
```

- OR -

Internal Only - General

```
# apt dist-upgrade
```

Note: When running the command `apt dist-upgrade` that apt has a "smart" conflict resolution system, and it will attempt to upgrade the most important packages at the expense of less important ones if necessary. So, `dist-upgrade` command may remove some packages.

Once the update process is complete, verify if a reboot is required to load changes:

```
# if [ -f /var/run/reboot-required ]; then
     cat /var/run/reboot-required
  fi
```

References:

1. GRID: NIX-00000081
2. NIST SP 800-53 Rev. 5: SI-2

Additional Information:

Site policy may mandate a testing period before installation onto production systems for available updates.

- upgrade - is used to install the newest versions of all packages currently installed on the system from the sources enumerated in /etc/apt/sources.list. Packages currently installed with new versions available are retrieved and upgraded; under no circumstances are currently installed packages removed, or packages not already installed retrieved and installed. New versions of currently installed packages that cannot be upgraded without changing the install status of another package will be left at their current version. An update must be performed first so that apt knows that new versions of packages are available.
- dist-upgrade - in addition to performing the function of upgrade, also intelligently handles changing dependencies with new versions of packages; apt has a "smart" conflict resolution system, and it will attempt to upgrade the most important packages at the expense of less important ones if necessary. So, dist-upgrade command may remove some packages. The /etc/apt/sources.list file contains a list of locations from which to retrieve desired package files. See also apt_preferences(5) for a mechanism for overriding the general settings for individual packages.

Internal Only - General

CIS Controls:

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

1.3 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.

Internal Only - General

1.3.1 Configure AppArmor

AppArmor provides a Mandatory Access Control (MAC) system that greatly augments the default Discretionary Access Control (DAC) model. Under AppArmor MAC rules are applied by file paths instead of by security contexts as in other MAC systems. As such it does not require support in the filesystem and can be applied to network mounted filesystems for example. AppArmor security policies define what system resources applications can access and what privileges they can do so with. 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 AppArmor MAC rules. The action will not be allowed if either one of these models does not permit the action. In this way, AppArmor rules can only make a system's permissions more restrictive and secure.

References:

1. AppArmor Documentation: http://wiki.apparmor.net/index.php/Documentation
2. Ubuntu AppArmor Documentation: https://help.ubuntu.com/community/AppArmor
3. SUSE AppArmor Documentation: https://www.suse.com/documentation/apparmor/

Internal Only - General

1.3.1.1 Ensure apparmor packages are installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

AppArmor provides Mandatory Access Controls.

Rationale:

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

Audit:

Run the following command to verify that `apparmor` is installed:

```
# dpkg-query -s apparmor &>/dev/null && echo "apparmor is installed"

apparmor is installed
```

Run the following command to verify that `apparmor-utils` is installed:

```
# dpkg-query -s apparmor-utils &>/dev/null && echo "apparmor-utils is
installed"

apparmor-utils is installed
```

Remediation:

Run the following command to install `apparmor` and `apparmor-utils`:

```
# apt install apparmor apparmor-utils
```

References:

1. GRID: NIX-00000087
2. NIST SP 800-53 Rev. 5: AC-3

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.	●	●	●

Internal Only - General

Controls Version	Control	IG 1	IG 2	IG 3
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.	●	●	●

Internal Only - General

1.3.1.2 Ensure AppArmor is enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`AppArmor` is a kernel enhancement to confine programs to a limited set of resources. `AppArmor` is enabled by default.

Rationale:

`AppArmor` is a security mechanism and disabling it is not recommended.

Audit:

Run the following commands to verify that `AppArmor` has not been disabled at the kernel command line:

```
# grep "^\s*linux" /boot/grub/grub.cfg | grep "apparmor=0"
# grep "apparmor=0" /etc/default/grub.d/apparmor.cfg 2>/dev/null
```

Nothing should be returned from either command.

Run the following commands to verify `AppArmor` is enabled and active at runtime:

```
# systemctl is-enabled apparmor.service 2>/dev/null
```

Expected output: `enabled`

```
# systemctl is-active apparmor.service 2>/dev/null
```

Expected output: `active`

```
# /usr/sbin/apparmor_status 2>/dev/null | grep 'profiles are loaded'
```

Expected output: `<N> profiles are loaded` where N is greater than 0.

If `apparmor.service` is not enabled and active, or if no profiles are loaded, this recommendation is not in compliance.

Remediation:

Edit `/etc/default/grub` or a file in `/etc/default/grub.d` and remove the `apparmor=0` parameter from the `GRUB_CMDLINE_LINUX=` line.

Run the following commands to update the grub configuration and reboot the system:

Internal Only - General

```
# update-grub
# reboot
```

References:

1. GRID: NIX-00000085
2. NIST SP 800-53 Rev. 5: AC-3

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.	●	●	●

Internal Only - General

1.3.1.3 Ensure all AppArmor Profiles are enforcing (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

AppArmor profiles define what resources applications are able to access.

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 any policies that exist on the system are enforcing, actively denying policy violations rather than merely logging them.

Impact:

Setting AppArmor profiles to `enforce` mode causes policy violations to be denied rather than logged. Applications with incomplete or overly-restrictive profiles may fail or behave unexpectedly. Profiles should be tested in `complain` mode and tuned before being set to `enforce`.

> **Note:** AppArmor 4, shipped with Debian 13, introduces "stub profiles" located in `/etc/apparmor.d/` (for example, `busybox`). Stub profiles contain only a `userns,` rule they assign a label to a process for user namespace tracking but intentionally have no file, capability, or network access rules. Enforcing a stub profile applies AppArmor's default-deny policy, causing all file access to be denied and the associated application to fail. All stub profiles carry the identifying comment: `"This profile exist only to give a name rather than the label unconfined"`.
>
> Stub profiles must be fully developed with appropriate file and capability rules before being set to enforce mode. See Additional Information for guidance on identifying and tuning stub profiles.

Audit:

Run the following command and verify the output:

```
# /usr/sbin/apparmor_status | grep profiles
```

Example Output

Internal Only - General

```
34 profiles are loaded.
34 profiles are in enforce mode.
0 profiles are in complain mode.
0 profiles are in prompt mode.
0 profiles are in kill mode.
0 profiles are in unconfined mode.
0 processes have profiles defined.
```

> **Note:** The count of profiles in enforce mode must equal the total profiles loaded. The count of profiles in all other modes (complain, prompt, kill, unconfined) must be 0.

If the system returns a non-zero count for profiles in complain mode, run the following to identify whether any are stub profiles requiring tuning before enforcement:

```
# grep -rl "This profile exist only to give a name" /etc/apparmor.d/
```

Any profile returned is a stub profile that must be developed with appropriate access rules and enforced before the system meets the Level 2 requirement. See Additional Information for guidance.

Run the following command and verify the output:

```
# /usr/sbin/apparmor_status | grep processes
```

Example Output

```
2 processes have profiles defined.
2 processes are in enforce mode.
0 processes are in complain mode.
0 processes are in prompt mode.
0 processes are in kill mode.
0 processes are unconfined but have a profile defined.
0 processes are in mixed mode.
```

> **Note:** The count of processes unconfined but have a profile defined must be 0.

Remediation:

First, identify any stub profiles present on the system:

```
# grep -rl "This profile exist only to give a name" /etc/apparmor.d/
```

For each stub profile returned, either:

- Develop the profile by adding appropriate file, capability, and network rules for the application, then enforce it individually:

```
# /usr/sbin/aa-enforce /etc/apparmor.d/<profile-name>
```

- Or remove the stub if the application is not present or needed on the system:

Internal Only - General

```
# apt remove <package-name>
```

Once all stub profiles have been developed and enforced, or removed, run the following to enforce all remaining profiles:

```
# /usr/sbin/aa-enforce /etc/apparmor.d/*
```

> **Note:** The glob `/etc/apparmor.d/*` matches top-level profile files as well as subdirectories (`abstractions/`, `tunables/`, `abi/`, `local/`). Passing a subdirectory to `aa-enforce` may produce warnings but does not affect profile enforcement. Any unconfined processes may need to have a profile created or activated for them and then be restarted.

References:

1. GRID: NIX-00000083
2. NIST SP 800-53 Rev. 5: AC-3

Additional Information:

Identifying stub profiles (AppArmor 4 / Debian 13)

To list all stub profiles on the system:

```
# grep -rl "This profile exist only to give a name" /etc/apparmor.d/
```

Stub profiles are shipped in complain mode by design. They contain only a `userns,` rule to assign a process label for user namespace tracking. Because AppArmor operates under a default-deny model, enforcing a stub profile without first adding file and capability rules will deny all file access for that process.

Developing a stub profile

Use AppArmor's log-based profile generator to interactively build the necessary rules:

```
# aa-logprof /etc/apparmor.d/<profile-name>
```

This tool reads AppArmor audit events from the system log and proposes rules to add. After adding rules, test the application in complain mode before enforcing:

```
# /usr/sbin/aa-complain /etc/apparmor.d/<profile-name>
# <run application workload>
# /usr/sbin/aa-logprof /etc/apparmor.d/<profile-name>
# /usr/sbin/aa-enforce /etc/apparmor.d/<profile-name>
```

Selective enforcement during profile development (not a compliant state)

During the development and testing phase, the following may be used to enforce only non-stub profiles while leaving stub profiles in complain mode until they are tuned. A system in this intermediate state does not meet the Level 2 requirement:

Internal Only - General

```
for f in /etc/apparmor.d/*; do
  [ -f "$f" ] || continue
  grep -q "This profile exist only to give a name" "$f" 2>/dev/null &&
continue
  /usr/sbin/aa-enforce "$f"
done
```

Reference

https://gitlab.com/apparmor/apparmor/-/wikis/AppArmorUserNamespaces

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.	●	●	●

Internal Only - General

1.3.1.4 Ensure kernel.apparmor_restrict_unprivileged_unconfined is enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

By making sure the `kernel.apparmor_restrict_unprivileged_unconfined` sysctl setting is enabled, a process that is unprivileged and unconfined cannot leverage `aa-exec` to change to a more favourable profile

Rationale:

Unprivileged user namespaces are a feature in the Linux kernel that was introduced in order to provide additional sandboxing functionality for programs such as container runtimes; it enables unprivileged users to gain administrator (root) permissions within a confined environment, without giving them elevated permissions on the host system. However, unconfined profiles that allow the use of user namespaces provide a trivial bypass of the unprivileged user namespace restriction from unconfined processes. All the unconfined application or user has to do is enter the unconfined profile. The unconfined user or application can execute the application associated with the unconfined profile. If the application provides the ability to run arbitrary code or shell access, that is all that needs to be done to bypass the restriction

Audit:

1. Run the following script to verify `kernel.apparmor_restrict_unprivileged_unconfined` is set to 1 in the running configuration:

```
#  sysctl kernel.apparmor_restrict_unprivileged_unconfined
```

Verify output is:

```
kernel.apparmor_restrict_unprivileged_unconfined = 1
```

2. Run the following script to verify `kernel.apparmor_restrict_unprivileged_unconfined` is set to `1`in a sysctl conf file being used by systemd sysctl:

Internal Only - General

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

{
   l_parameter_name="kernel.apparmor_restrict_unprivileged_unconfined"
l_grep="${l_parameter_name//./\\.}" a_output=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl || readlink -e
/usr/lib/systemd/systemd-sysctl)"
   l_ufwscf="$([ -f /etc/default/ufw ] && awk -F= '/^\s*IPT_SYSCTL=/ {print
$2}' /etc/default/ufw)"
   l_opt="$(grep -Psoi '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_ufwscf" | tail -n
1)"
   l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
   [ -n "$l_option_value" ] && a_output+=(" - UFW set: \"$l_parameter_name\"
to: \"$l_option_value\" in: \"$l_file\"")
   while IFS= read -r l_file; do
      l_file="${l_file//# /}"
      l_opt="$(grep -Psoi '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && a_output+=(" - \"$l_parameter_name\" is set
to: \"$l_option_value\" in: \"$l_file\"")
   done < <("$l_systemdsysctl" --cat-config | tac | grep -Psoi
'^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Example Output

```
"kernel.apparmor_restrict_unprivileged_unconfined" is set to: "1" in:
"/etc/sysctl.d/60-kernel_sysctl.conf"
```

Note:

- If the UFW set value is displayed, and is the correct value, this is considered a passing state. If this value is incorrect, it must be updated in the file shown in the audit script output.
- If multiple lines are returned:
  - The first line includes the value being used by systemd `sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a .conf file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
  - Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to correct or comment out the incorrect value to minimize the potential of the incorrect value being used by systemd `sysctl` do to system configuration changes.

Internal Only - General

Remediation:

1. Set the `kernel.apparmor_restrict_unprivileged_unconfined =` kernel parameter to `1` in `/etc/sysctl.conf` or a file in `/etc/sysctl.d/` ending in `.conf`:

Example:

```
# printf "\n%s\n" "kernel.apparmor_restrict_unprivileged_unconfined = 1" >>
/etc/sysctl.d/60-kernel_sysctl.conf
```

2. Run the following command to set the active kernel parameter:

```
# sysctl -w kernel.apparmor_restrict_unprivileged_unconfined=1
```

3. Run the following command to load all system configuration filles:

```
# sysctl --system
```

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

References:

1. GRID: NIX-00000088
2. https://discourse.ubuntu.com/t/understanding-apparmor-user-namespace-restriction/58007

CIS Controls:

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

Internal Only - General

1.4 Configure Bootloader

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

Internal Only - General

1.4.1 Ensure bootloader 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 AppArmor 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 to do so, the configuration files will have to be edited via a 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. A password will still be required to edit menu items.

More Information: https://help.ubuntu.com/community/Grub2/Passwords

Audit:

Run the following commands and verify output matches:

```
# grep "^set superusers" /boot/grub/grub.cfg

set superusers="<username>"
# awk -F. '/^\s*password/ {print $1"."$2"."$3}' /boot/grub/grub.cfg

password_pbkdf2 <username> grub.pbkdf2.sha512
```

Remediation:

Create an encrypted password with `grub-mkpasswd-pbkdf2`:

Internal Only - General

```
# grub-mkpasswd-pbkdf2 --iteration-count=600000 --salt=64

Enter password: <password>
Reenter password: <password>
PBKDF2 hash of your password is <encrypted-password>
```

Add the following into a custom `/etc/grub.d` configuration file:

```
set superusers="<username>"
password_pbkdf2 <username> <encrypted-password>
```

The superuser/user information and password should not be contained in the `/etc/grub.d/00_header` file as this file could be overwritten in a package update.

If there is a requirement to be able to boot/reboot without entering the password, edit `/etc/grub.d/10_linux` and add `--unrestricted` to the line `CLASS=`

Example:

```
CLASS="--class gnu-linux --class gnu --class os --unrestricted"
```

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

```
# update-grub
```

Default Value:

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

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

References:

1. GRID: NIX-00000098
2. NIST SP 800-53 Rev. 5: AC-3
3. NIST SP 800-53A :: AC-3.1

Additional Information:

Changes to `/etc/grub.d/10_linux` may be overwritten during updates to the `grub-common` package. You should review any changes to this file before rebooting otherwise the system may unexpectedly prompt for a password on the next boot.

Internal Only - General

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.	●	●	●

Internal Only - General

1.4.2 Ensure access to bootloader config is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The grub configuration file contains information on boot settings and passwords for unlocking boot options.

Rationale:

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

Audit:

Run the following command and verify `Uid` and `Gid` are both `0/root` and `Access` is `0600` or more restrictive.

```
# stat -Lc 'Access: (%#a/%A)  Uid: ( %u/ %U) Gid: ( %g/ %G)'
/boot/grub/grub.cfg

Access: (0600/-rw-------)  Uid: ( 0/ root) Gid: ( 0/ root)
```

Remediation:

Run the following commands to set permissions on your grub configuration:

```
# chown root:root /boot/grub/grub.cfg
# chmod u-x,go-rwx /boot/grub/grub.cfg
```

Default Value:

Access: (0644/-rw-r--r--) Uid: (0/ root) Gid: (0/ root)

References:

1. GRID: NIX-00000097
2. NIST SP 800-53 Rev. 5: AC-3

Additional Information:

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

Internal Only - General

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

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.	●	●	●

Internal Only - General

1.5 Configure Additional Process Hardening

This section includes Recommendations covering kernel parameters.

- sysctl settings are defined through files in `/usr/local/lib`, `/usr/lib/`, `/lib/`, `/run/`, and `/etc/`
- Files are typically placed in the `sysctl.d` directory within the parent directory
- The paths where sysctl preload files usually exist
    - `/run/sysctl.d/*.conf`
    - `/etc/sysctl.d/*.conf`
    - `/usr/local/lib/sysctl.d/*.conf`
    - `/usr/lib/sysctl.d/*.conf`
    - `/lib/sysctl.d/*.conf`
    - `/etc/sysctl.conf`
- Files must have the "`.conf`" extension
- Vendors settings usually live in `/usr/lib/` or `/usr/local/lib/`
- To override a whole file, create a new file with the same name in `/etc/sysctl.d/` and put new settings there.
- To override only specific settings, add a file with a lexically later name in `/etc/sysctl.d/` and put new settings there.
- The command `/usr/lib/systemd/systemd-sysctl --cat-config` produces output containing The system's loaded kernel parameters and the files they're configured in:
    - Entries listed latter in the file take precedence over the same settings listed earlier in the file
    - Files containing kernel parameters that are over-ridden by other files with the same name will not be listed
    - On systems running UncomplicatedFirewall, the kernel parameters may be set or over-written. This will not be visible in the output of the command
- On systems with Uncomplicated Firewall, additional settings may be configured in `/etc/ufw/sysctl.conf`
- The settings in `/etc/ufw/sysctl.conf` will override settings other settings and **will not** be visible in the output of the `/usr/lib/systemd/systemd-sysctl --cat-config` command
- This behavior can be changed by updating the `IPT_SYSCTL` parameter in `/etc/default/ufw`
- **SYSTEM FILE PRECEDENCE**
    - When using the `--system` option, `sysctl` will read files from directories in the following list in given order from top to bottom. Once a file of a given filename is loaded, any file of the same name in subsequent directories is ignored.

      ```
      /etc/sysctl.d/*.conf /run/sysctl.d/*.conf
      /usr/local/lib/sysctl.d/*.conf /usr/lib/sysctl.d/*.conf
      /lib/sysctl.d/*.conf /etc/sysctl.conf
      ```


Internal Only - General

- All configuration files are sorted in lexicographic order, regardless of the directory they reside in. Configuration files can either be completely replaced (by having a new configuration file with the same name in a directory of higher priority) or partially replaced (by having a configuration file that is ordered later)--

Internal Only - General

1.5.1 Ensure fs.protected_hardlinks is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`fs.protected_hardlinks` Controls the creation of hard links to files.

More information about the kernel parameter configuration files, their location, and load preference is available in the section overview.

Rationale:

Disallowing hardlinks mitigates vulnerabilities based on unsecure file systems accessed by privileged programs. This reduces the risk of an exploitation vector exploiting unsafe use of `open` or `creat`.

Impact:

Users will not be able to create soft or hard links to files they do not own.

Audit:

Verify `fs.protected_hardlinks` is set to `1`.

1. Run the following command to verify `fs.protected_hardlinks` is set to `1` in the running configuration:

```
# sysctl fs.protected_hardlinks
```

Verify output is:

```
fs.protected_hardlinks = 1
```

2. Run the following script to verify `fs.protected_hardlinks = 1` is set in a file being used by systemd-sysctl to configure `fs.protected_hardlinks`:

Internal Only - General

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

{
   l_parameter_name="fs.protected_hardlinks"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psoi '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:" \
      "    \"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `fs.protected_hardlinks = 1`.

Example output:

```
 - "fs.protected_hardlinks = 1" is set in: "/etc/sysctl.d/60-fs_sysctl.conf"
```

Verify output includes `fs.protected_hardlinks = 1`.

Note: If multiple lines with `fs.protected_hardlinks =` are returned:

- The first line includes the value being used by `systemd-sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect value to minimize the potential of the incorrect value being used by `systemd-sysctl` due to system configuration changes.

Remediation:

1. Review all files being used by systemd-sysctl and comment out or remove all `fs.protected_hardlinks` lines that are not `fs.protected_hardlinks=1`.

Internal Only - General

Example script:

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

{
   l_option="fs.protected_hardlinks" l_value="1"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(0|[2-9]|1[0-9]+)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf` and edit or add the following line:

```
fs.protected_hardlinks = 1
```

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "fs.protected_hardlinks = 1" >> /etc/sysctl.d/60-
fs_sysctl.conf
```

Note: If the UFW file was the first file listed in the audit, the entry will be commented out as part of the first step, however updating Uncomplicated Firewall (UFW) may update this change. In this case the updated entry will supersede the entry being created as part of this step.

3. Run the following command to load all sysctl configuration files:

```
# sysctl --system
```

Default Value:

fs.protected_hardlinks = 0

References:

1. GRID: NIX-00000101
2. NIST SP 800-53 :: AC-3 (4) | NIST SP 800-53 :: AC-6 (10)

Internal Only - General

3. STIG ID: RHEL-08-010374 | RULE ID: SV-230268r1017086 | CAT II
4. STIG ID: OL08-00-010374 | RULE ID: SV-248578r958702 | CAT II
5. STIG ID: RHEL-09-213030 | RULE ID: SV-257801r958702 | CAT II
6. STIG ID: ALMA-09-007060 | RULE ID: SV-269144r1050026 | CAT II
7. STIG ID: OL09-00-002401 | RULE ID: SV-271740r1091932 | CAT II
8. CCI-002165 | CCI-002235

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

1.5.2 Ensure fs.protected_symlinks is configured (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

`fs.protected_symlinks` Controls how the kernel handles symbolic links.

More information about the kernel parameter configuration files, their location, and load preference is available in the section overview.

Rationale:

Disallowing symlinks mitigates vulnerabilities based on unsecure file systems accessed by privileged programs. This reduces the risk of an exploitation vector exploiting unsafe use of `open` or `creat`.

Impact:

By enabling the `fs.protected_symlinks` kernel parameter, symbolic links are permitted to be followed only when outside a sticky world-writable directory, or when the user identifier (UID) of the link and follower match, or when the directory owner matches the symlink's owner.

Audit:

Verify `fs.protected_symlinks` is set to `1`.

1. Run the following command to verify `fs.protected_symlinks` is set to `1` in the running configuration:

```
# sysctl fs.protected_symlinks
```

Verify output is:

```
fs.protected_symlinks = 1
```

2. Run the following script to verify `fs.protected_symlinks = 1` is set in a file being used by systemd-sysctl to configure `fs.protected_symlinks`:

Internal Only - General

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

{
   l_parameter_name="fs.protected_symlinks"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psoi '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:" \
      "    \"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `fs.protected_symlinks = 1`.

Example output:

```
 - "fs.protected_symlinks = 1" is set in: "/etc/sysctl.d/60-fs_sysctl.conf"
```

Note: If multiple lines with `fs.protected_symlinks =` are returned:

- The first line includes the value being used by `systemd-sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect value to minimize the potential of the incorrect value being used by `systemd-sysctl` due to system configuration changes.

Remediation:

1. Review all files being used by systemd-sysctl and comment out or remove all `fs.protected_symlinks` lines that are not `fs.protected_symlinks=1`.

Example script:

Internal Only - General

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

{
   l_option="fs.protected_symlinks" l_value="1"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(0|[2-9]|1[0-9]+)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf` and edit or add the following line:

```
fs.protected_symlinks = 1
```

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "fs.protected_symlinks = 1" >> /etc/sysctl.d/60-
fs_sysctl.conf
```

Note: If the UFW file was the first file listed in the audit, the entry will be commented out as part of the first step, however updating Uncomplicated Firewall (UFW) may update this change. In this case the updated entry will supersede the entry being created as part of this step.

3. Run the following command to load all sysctl configuration files:

```
# sysctl --system
```

Default Value:

fs.protected_symlinks = 0

References:

1. GRID: NIX-00000102
2. NIST SP 800-53 :: AC-3 (4) | NIST SP 800-53 :: AC-6 (10)
3. STIG ID: RHEL-08-010374 | RULE ID: SV-230268r1017086 | CAT II

Internal Only - General

4. STIG ID: OL08-00-010373 | RULE ID: SV-248577r958702 | CAT II
5. STIG ID: RHEL-09-213030 | RULE ID: SV-257801r958702 | CAT II
6. STIG ID: ALMA-09-007060 | RULE ID: SV-269144r1050026 | CAT II
7. STIG ID: OL09-00-002402 | RULE ID: SV-271741r1091935 | CAT II
8. CCI-002165 | CCI-002235

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

1.5.3 Ensure kernel.yama.ptrace_scope is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

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

The sysctl settings (writable only with CAP_SYS_PTRACE) are:

- `0` - classic ptrace permissions: a process can PTRACE_ATTACH to any other process running under the same uid, as long as it is dumpable (i.e. did not transition uids, start privileged, or have called prctl(PR_SET_DUMPABLE...) already). Similarly, PTRACE_TRACEME is unchanged.
- `1` - restricted ptrace: a process must have a predefined relationship with the inferior it wants to call PTRACE_ATTACH on. By default, this relationship is that of only its descendants when the above classic criteria is also met. To change the relationship, an inferior can call prctl(PR_SET_PTRACER, debugger, ...) to declare an allowed debugger PID to call PTRACE_ATTACH on the inferior. Using PTRACE_TRACEME is unchanged.
- `2` - admin-only attach: only processes with CAP_SYS_PTRACE may use ptrace with PTRACE_ATTACH, or through children calling PTRACE_TRACEME.
- `3` - no attach: no processes may use ptrace with PTRACE_ATTACH nor via PTRACE_TRACEME. Once set, this sysctl value cannot be changed.

  Note: A logic bug in the Linux kernel's `__ptrace_may_access()` function (CVE-2026-46333) allows an unprivileged local user to gain ptrace access to processes they should not be able to reach. Setting `kernel.yama.ptrace_scope = 2` restricts ptrace to processes holding `CAP_SYS_PTRACE` only, cutting off the exploit path on unpatched systems.

Rationale:

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

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

Internal Only - General

Audit:

Verify `kernel.yama.ptrace_scope` is set to a value of: `1`, `2`, or `3`.

1. Run the following command to verify `kernel.yama.ptrace_scope` is set to a value of: `1`, `2`, or `3` in the running configuration:

```
# sysctl kernel.yama.ptrace_scope
```

Example output:

```
kernel.yama.ptrace_scope = 1
```

2. Run the following script to verify `kernel.yama.ptrace_scope` is set to a value of: `1`, `2`, or `3` in a sysctl conf file being used by systemd-sysctl:

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

{
   l_parameter_name="kernel.yama.ptrace_scope"
l_grep="${l_parameter_name//./\\.}" a_output=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl || readlink -e
/usr/lib/systemd/systemd-sysctl)"
   l_ufwscf="$([ -f /etc/default/ufw ] && awk -F= '/^\s*IPT_SYSCTL=/ {print
$2}' /etc/default/ufw)"
   l_opt="$(grep -Psoi '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_ufwscf" | tail -n
1)"
   l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
   [ -n "$l_option_value" ] && a_output+=(" - UFW set: \"$l_parameter_name\"
to: \"$l_option_value\" in: \"$l_file\"")
   while IFS= read -r l_file; do
      l_file="${l_file//# /}"
      l_opt="$(grep -Pois '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && a_output+=(" - \"$l_parameter_name\" is set
to: \"$l_option_value\" in: \"$l_file\"")
   done < <("$l_systemdsysctl" --cat-config | tac | grep -Pios
'^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Example output:

```
- "kernel.yama.ptrace_scope" is set to: "1" in: "/etc/sysctl.d/60-
kernel_sysctl.conf"
```

Note:

- If the `UFW set` value is displayed, and is the correct value, this is considered a passing state. If this value is incorrect, it must be updated in the file shown in the audit script output.
- If multiple lines are returned:

Internal Only - General

- The first line includes the value being used by `systemd-sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to correct or comment out the incorrect value to minimize the potential of the incorrect value being used by `systemd-sysctl` do to system configuration changes.

Remediation:

1. Run the following script to comment out any non-compliant `kernel.yama.ptrace_scope` entries from all sysctl configuration files enumerated by systemd-sysctl:

Example script:

```
#!/usr/bin/env bash
{
   l_option="kernel.yama.ptrace_scope" l_value="(1|2|3)"
   l_grep="${l_option//./(\\.|\\/)}"; a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Pios '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Pivq '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(0|[4-9]|[1-9][0-9]+)/s/^/# /'
"$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

```
kernel.yama.ptrace_scope = 1
```

Example:

Internal Only - General

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "kernel.yama.ptrace_scope = 1" >> /etc/sysctl.d/60-
kernel_sysctl.conf
```

Note:

- The example uses `kernel.yama.ptrace_scope = 1` but value may be set to `1`, `2`, or `3`
- A logic bug in the Linux kernel's `__ptrace_may_access()` function (CVE-2026-46333) allows an unprivileged local user to gain ptrace access to processes they should not be able to reach. Setting `kernel.yama.ptrace_scope = 2` restricts ptrace to processes holding `CAP_SYS_PTRACE` only, cutting off the exploit path on unpatched systems.

3. Run the following command to load all system configuration files:

```
# sysctl --system
```

Note: If the UFW file was the first file listed in the audit, the entry will be commented out as part of the first step, however updating Uncomplicated Firewall (UFW) may update this change. In this case the updated entry will supersede the entry being created as part of this step.

Default Value:

kernel.yama.ptrace_scope = 0

References:

1. GRID: NIX-00000107
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-08-040282 | RULE ID: SV-230546r1017308 | CAT II
5. STIG ID: RHEL-09-213080 | RULE ID: SV-257811r1044872 | CAT II
6. STIG ID: ALMA-09-040940 | RULE ID: SV-269424r1050307 | CAT II
7. CCI-000366

Additional Information:

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

Internal Only - General

CIS Controls:

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

Internal Only - General

1.5.4 Ensure fs.suid_dumpable is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`fs.suid_dumpable` governs whether a privileged process (with the setuid bit) can generate a core dump, regardless of other configurations.

`fs.suid_dumpable` values:

- `0` (default) - Any process that has changed privilege levels (like SUID programs) or is execute-only will not dump core.
- `1` (debug) - All processes dump core if possible. The core dump is owned by the current user and security is not applied. This is primarily intended for system debugging.
- `2` (suidsafe) - Any binary that normally wouldn't be dumped is dumped, but only if the core_pattern is set to a pipe handler or a fully qualified path. This mode is suitable for administrators debugging in a production environment.

More information about the kernel parameter configuration files, their location, and load preference is available in the section overview.

Rationale:

core dumps may contain sensitive in-memory data like password hashes or keys. An attacker could potentially exploit this to gain access to such data.

Impact:

Any process that has changed privilege levels (like SUID programs) or is execute-only will not dump core.

Audit:

Verify `fs.suid_dumpable` is set to `0`.

1. Run the following command to verify `fs.suid_dumpable` is set to `0` in the running configuration:

```
# sysctl fs.suid_dumpable
```

Verify output is:

Internal Only - General

```
fs.suid_dumpable = 0
```

2. Run the following script to verify `fs.suid_dumpable = 0` is set in a file being used by systemd-sysctl to configure `fs.suid_dumpable`:

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

{
   l_parameter_name="fs.suid_dumpable"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psoi '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:" \
      "    \"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `fs.suid_dumpable = 0`.

Example output:

```
 - "fs.suid_dumpable = 0" is set in: "/etc/sysctl.d/60-fs_sysctl.conf"
```

Note: If multiple lines with `fs.suid_dumpable =` are returned:

- The first line includes the value being used by `systemd-sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect value to minimize the potential of the incorrect value being used by `systemd-sysctl` due to system configuration changes.

Internal Only - General

Remediation:

1. Review all files being used by systemd-sysctl and comment out or remove all `fs.suid_dumpable` lines that are not `fs.suid_dumpable=0`.

Example script:

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

{
   l_option="fs.suid_dumpable" l_value="0"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf` and edit or add the following line:

```
fs.suid_dumpable = 0
```

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "fs.suid_dumpable = 0" >> /etc/sysctl.d/60-fs_sysctl.conf
```

Note: If the UFW file was the first file listed in the audit, the entry will be commented out as part of the first step, however updating Uncomplicated Firewall (UFW) may update this change. In this case the updated entry will supersede the entry being created as part of this step.

3. Run the following command to load all sysctl configuration files:

```
# sysctl --system
```

Default Value:

fs.suid_dumpable = 0

Internal Only - General

References:

1. GRID: NIX-00000103
2. NIST SP 800-53 Revision 5::AC-3 (4)

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

1.5.5 Ensure kernel.dmesg_restrict is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`kernel.dmesg_restrict` can be used to configure the Linux kernel and prevent unprivileged users access to information from the kernel's log buffer.

More information about the kernel parameter configuration files, their location, and load preference is available in the section overview.

Rationale:

Restricting access to the kernel message buffer limits access to only root. This prevents attackers from gaining additional system information as a non-privileged user.

Impact:

Changing dmesg_restrict to one `1`, will restrict access to those users that have the CAP_SYSLOG capability.

Audit:

Verify `kernel.dmesg_restrict` is set to `1`.

1. Run the following command to verify `kernel.dmesg_restrict` is set to `1` in the running configuration:

```
# sysctl kernel.dmesg_restrict
```

Verify output is:

```
kernel.dmesg_restrict = 1
```

2. Run the following script to verify `kernel.dmesg_restrict = 1` is set in a file being used by systemd-sysctl to configure `kernel.dmesg_restrict`:

Internal Only - General

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

{
   l_parameter_name="kernel.dmesg_restrict"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psoi '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:" \
      "    \"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `kernel.dmesg_restrict = 1`.

Example output:

```
 - "kernel.dmesg_restrict" is set to: "1" in: "/etc/sysctl.d/60-
kernel_sysctl.conf"
```

Note: If multiple lines with `kernel.dmesg_restrict =` are returned:

- The first line includes the value being used by `systemd-sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect value to minimize the potential of the incorrect value being used by `systemd-sysctl` due to system configuration changes.

Remediation:

1. Review all files being used by systemd-sysctl and comment out or remove all `kernel.dmesg_restrict` lines that are not `kernel.dmesg_restrict=1`.

Internal Only - General

Example script:

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

{
   l_option="kernel.dmesg_restrict" l_value="1"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(0|[2-9]|1[0-9]+)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf` and edit or add the following line:

```
kernel.dmesg_restrict = 1
```

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "kernel.dmesg_restrict = 1" >> /etc/sysctl.d/60-
kernel_sysctl.conf
```

Note: If the UFW file was the first file listed in the audit, the entry will be commented out as part of the first step, however updating Uncomplicated Firewall (UFW) may update this change. In this case the updated entry will supersede the entry being created as part of this step.

3. Run the following command to load all system configuration files:

```
# sysctl --system
```

Default Value:

kernel.dmesg_restrict = 0

References:

1. GRID: NIX-00000104
2. NIST SP 800-53 :: SC-2 | NIST SP 800-53 :: SC-4

Internal Only - General

3. NIST SP 800-53A :: SC-2.1 | NIST SP 800-53A :: SC-4.1
4. STIG ID: RHEL-08-010375 | RULE ID: SV-230269r1017087 | CAT III
5. STIG ID: OL08-00-010375 | RULE ID: SV-248579r958524 | CAT III
6. STIG ID: RHEL-09-213010 | RULE ID: SV-257797r958514 | CAT II
7. STIG ID: ALMA-09-041050 | RULE ID: SV-269425r1050308 | CAT II
8. STIG ID: OL09-00-002406 | RULE ID: SV-271745r1091947 | CAT II
9. STIG ID: UBTU-22-213010 | RULE ID: SV-260472r958524 | CAT III
10. STIG ID: UBTU-24-600140 | RULE ID: SV-270749r1067179 | CAT III
11. STIG ID: SLES-15-010375 | RULE ID: SV-255921r958524 | CAT III
12. CCI-001082 | CCI-001090

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.	●	●	●

Internal Only - General

1.5.6 Ensure prelink is not installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`prelink` is a program that modifies ELF shared libraries and ELF dynamically linked binaries in such a way that the time needed for the dynamic linker to perform relocations at startup significantly decreases.

Rationale:

The prelinking feature can interfere with the operation of AIDE, because it changes binaries. Prelinking can also increase the vulnerability of the system if a malicious user is able to compromise a common library such as libc.

Audit:

Verify `prelink` is not installed:

```
# dpkg-query -s prelink &>/dev/null && echo "prelink is installed"
```

Nothing should be returned.

Remediation:

Run the following command to restore binaries to normal:

```
# prelink -ua
```

Uninstall `prelink` using the appropriate package manager or manual installation:

```
# apt purge prelink
```

References:

1. GRID: NIX-00000108
2. NIST SP 800-53 Rev. 5: CM-6, CM-1, CM-3

CIS Controls:

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

Internal Only - General

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

Internal Only - General

1.5.7 Ensure Automatic Error Reporting is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The Apport Error Reporting Service automatically generates crash reports for debugging

Rationale:

Apport collects potentially sensitive data, such as core dumps, stack traces, and log files. They can contain passwords, credit card numbers, serial numbers, and other private material.

Audit:

Run the following command to verify that the Apport Error Reporting Service is not enabled:

```
#  dpkg-query -s apport &> /dev/null && grep -Psi --
'^\h*enabled\h*=\h*[^0]\b' /etc/default/apport
```

Nothing should be returned

Run the following command to verify that the apport service is not active:

```
# systemctl is-active apport.service | grep '^active'
```

Nothing should be returned

Remediation:

Edit `/etc/default/apport` and add or edit the enabled parameter to equal `0`:

```
enabled=0
```

Run the following commands to stop and mask the apport service

```
# systemctl stop apport.service
# systemctl mask apport.service
```

- OR -

Run the following command to remove the apport package:

```
# apt purge apport
```

Default Value:

enabled=1

Internal Only - General

References:

1. GRID: NIX-00000099

CIS Controls:

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

Internal Only - General

1.5.8 Ensure kernel.kptr_restrict is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`kernel.kptr_restrict` sets whether restrictions are placed on exposing kernel addresses via /proc and other interfaces.

- `kptr_restrict = 0` - The address is hashed before printing. (equivalent to %p.)
- `kptr_restrict = 1` - kernel pointers printed using the %pK format specifier will be replaced with `0`'s.
- `kptr_restrict = 2` - kernel pointers printed using %pK will be replaced with `0`'s regardless of privileges.

More information about the kernel parameter configuration files, their location, and load preference is available in the section overview.

Rationale:

Masking the kernel symbols in `/proc/kallsyms` reduces the ability of an attacker using them to learn more about what to attack on your system.

Audit:

Verify `kernel.kptr_restrict` is set to a value of: `1` or `2`.

1. Run the following command to verify `kernel.kptr_restrict` is set to a value of: `1` or `2` in the running configuration:

```
# sysctl kernel.kptr_restrict
```

Example output:

```
kernel.kptr_restrict = 2
```

2. Run the following script to verify `kernel.kptr_restrict` is set to a value of: `1` or `2` in a sysctl conf file being used by systemd-sysctl:

Internal Only - General

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

{
   l_parameter_name="kernel.kptr_restrict"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psoi '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:" \
      "    \"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes one of the following:

- `kernel.kptr_restrict = 1`
- `kernel.kptr_restrict = 2`

Example output:

```
- "kernel.kptr_restrict" is set to: "1" in: "/etc/sysctl.d/60-
kernel_sysctl.conf"
```

Note: If multiple lines with `kernel.kptr_restrict =` are returned:

- The first line includes the value being used by `systemd-sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect value to minimize the potential of the incorrect value being used by `systemd-sysctl` due to system configuration changes.

Internal Only - General

Remediation:

1. Review all files being used by systemd-sysctl and comment out or remove all `kernel.kptr_restrict` lines that are not `kernel.kptr_restrict=1` or `kernel.kptr_restrict=2`.

Example script:

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

{
   l_option="kernel.kptr_restrict" l_value="(1|2)"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Pios '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Pios '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Pivq '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(0|[3-9]|[1-9][0-9]+)/s/^/# /'
"$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf` and edit or add the following line:

```
kernel.kptr_restrict = {N}
```

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "kernel.kptr_restrict = 2" >> /etc/sysctl.d/60-
kernel_sysctl.conf
```

Note: The example uses `kernel.kptr_restrict = 2` but value may be set to `1` or `2`

3. Run the following command to load all system configuration files:

Internal Only - General

```
# sysctl --system
```

Note: If the UFW file was the first file listed in the audit, the entry will be commented out as part of the first step, however updating Uncomplicated Firewall (UFW) may update this change. In this case the updated entry will supersede the entry being created as part of this step.

Default Value:

kernel.kptr_restrict = 1

References:

1. GRID: NIX-00000105
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-08-040283 | RULE ID: SV-230547r1017309 | CAT II
5. STIG ID: OL08-00-040283 | RULE ID: SV-248891r991589 | CAT II
6. STIG ID: RHEL-09-213025 | RULE ID: SV-257800r1044851 | CAT II
7. STIG ID: ALMA-09-040830 | RULE ID: SV-269423r1050306 | CAT II
8. STIG ID: OL09-00-002408 | RULE ID: SV-271747r1091953 | CAT II
9. STIG ID: SLES-15-010540 | RULE ID: SV-234861r958928 | CAT II
10. CCI-000366

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	4.8 Log and Alert on Changes to Administrative Group Membership Configure systems to issue a log entry and alert when an account is added to or removed from any group assigned administrative privileges.		●	●
v7	9.2 Ensure Only Approved Ports, Protocols and Services Are Running Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.		●	●

Internal Only - General

1.5.9 Ensure kernel.randomize_va_space is configured (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.

More information about the kernel parameter configuration files, their location, and load preference is available in the section overview.

Rationale:

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

Audit:

Verify `kernel.randomize_va_space` is set to `2`.

1. Run the following command to verify `kernel.randomize_va_space` is set to `2` in the running configuration:

```
# sysctl kernel.randomize_va_space
```

Verify output is:

```
kernel.randomize_va_space = 2
```

2. Run the following script to verify `kernel.randomize_va_space = 2` is set in a file being used by systemd sysctl to configure `kernel.randomize_va_space`:

Internal Only - General

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

{
   l_parameter_name="kernel.randomize_va_space"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psoi '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:" \
      "    \"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `kernel.randomize_va_space = 2`.

Example output:

```
 - "kernel.randomize_va_space" is set to: "2" in: "/etc/sysctl.d/60-
kernel_sysctl.conf"
```

Note: If multiple lines with `kernel.randomize_va_space =` are returned:

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `kernel.randomize_va_space` lines that are not `kernel.randomize_va_space=2`.

Internal Only - General

Example script:

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

{
   l_option="kernel.randomize_va_space" l_value="2"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*([01]|[3-9]|1[0-9]+)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf` and edit or add the following line:

```
kernel.randomize_va_space = 2
```

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "kernel.randomize_va_space = 2" >> /etc/sysctl.d/60-
kernel_sysctl.conf
```

Note: If the UFW file was the first file listed in the audit, the entry will be commented out as part of the first step, however updating Uncomplicated Firewall (UFW) may update this change. In this case the updated entry will supersede the entry being created as part of this step.

3. Run the following command to load all system configuration filles:

```
# sysctl --system
```

Default Value:

kernel.randomize_va_space = 2

References:

1. GRID: NIX-00000106
2. NIST SP 800-53 :: SI-16

Internal Only - General

3. STIG ID: RHEL-08-010430 | RULE ID: SV-230280r1017093 | CAT II
4. STIG ID: OL08-00-010430 | RULE ID: SV-248594r958928 | CAT II
5. STIG ID: RHEL-09-213070 | RULE ID: SV-257809r1044866 | CAT II
6. STIG ID: ALMA-09-044900 | RULE ID: SV-269452r1050335 | CAT II
7. STIG ID: OL09-00-002423 | RULE ID: SV-271761r1091995 | CAT II
8. STIG ID: UBTU-22-213020 | RULE ID: SV-260474r958928 | CAT II
9. STIG ID: UBTU-24-700310 | RULE ID: SV-270772r1066805 | CAT II
10. STIG ID: SLES-15-010550 | RULE ID: SV-234862r958928 | CAT II
11. CCI-002824

CIS Controls:

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

Internal Only - General

1.5.10 Ensure core file size is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

core - limits the core file size

hard - for enforcing hard resource limits. These limits are set by the superuser and enforced by the Kernel. The user cannot raise their requirement of system resources above such values.

Rationale:

Setting a hard limit on core dumps prevents users from overriding the soft variable.

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

Audit:

Run the following command to verify a hard limit for `core` is set to `0` for all users `*`:

```
# grep -Psi -- '^\h*\*\h+hard\h+core\b' /etc/security/limits.conf
/etc/security/limits.d/*
```

Example output:

```
/etc/security/limits.d/60-limits.conf:* hard core 0
```

Verify no line are returned with a value greater than `0`.

> **Note**: By default limits are taken from the `/etc/security/limits.conf` config file. Then individual `*.conf` files from the `/etc/security/limits.d/` directory are read. The files are parsed one after another in the order of "C" locale. The effect of the individual files is the same as if all the files were concatenated together in the order of parsing. When the same limit is defined in multiple files, the last matching entry takes precedence. If a config file is explicitly specified with a module option then the files in the above directory are not parsed.

Remediation:

1. Run the following command to comment out any entries that include a `hard` value for `core` greater than `0` in `/etc/security/limits.conf` and any file(s) in the `/etc/security/limits.d/` directory.

Internal Only - General

Example:

```
# sed -ri '/^\s*[^#\n\r]+\s+hard\s+core\s+([1-9][0-9]*)/s/^/# /'
/etc/security/limits.conf  /etc/security/limits.d/*
```

2. Create or edit a file in `/etc/security/limits.d/` and add the following line:

```
* hard core 0
```

Example:

```
# printf '%s\n' "" "* hard core 0" >> /etc/security/limits.d/60-limits.conf
```

Default Value:

None

References:

1. GRID: NIX-00000100
2. CCI-000366
3. NIST SP 800-53 :: CM-6 b
4. NIST SP 800-53A :: CM-6.1 (iv)
5. STIG ID: RHEL-08-010673 | RULE ID: SV-230313r1069304 | CAT II
6. STIG ID: RHEL-09-213095 | RULE ID: SV-257814r991589 | CAT II
7. STIG ID: ALMA-09-044900 | RULE ID: SV-269452r1050335 | CAT II
8. limits.conf(5)

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

1.5.11 Ensure systemd-coredump ProcessSizeMax is configured (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.

Rationale:

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

Impact:

Setting ProcessSizeMax=0 effectively disables processing of core dumps by systemd-coredump.

Audit:

Run the following script to verify `systemd-coredump ProcessSizeMax` is set to `0`:

Internal Only - General

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

{
   l_analyze_cmd="$(readlink -e /bin/systemd-analyze || \
   readlink -e /usr/bin/systemd-analyze)"
   l_conf_file="systemd/coredump.conf" l_block="Coredump"
   l_option="ProcessSizeMax" l_option_value="" a_output=()
   while IFS= read -r l_file; do
      l_file="${l_file//# /}"
      l_opt="$(awk '/\['"$l_block"'\]/{a=1;next}/\[/{a=0}a' "$l_file" \
      2>/dev/null | grep -Poi '^\h*'"$l_option"'\h*=\h*\H+\b' | tail -n 1)"
         l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_option\" is set to: \"$l_option_value\"" \
      "    in: \"$l_file\"")
   done < <("$l_analyze_cmd" cat-config "$l_conf_file" | tac | \
   grep -Pio '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   if [ "${#a_output[@]}" -le "0" ]; then
      l_file="$(readlink -e /etc/"$l_conf_file" || \
      readlink -e /usr/lib/"$l_conf_file")"
      if [ -e "$l_file" ]; then
         l_opt="$(awk '/\['"$l_block"'\]/{a=1;next}/\[/{a=0}a' "$l_file" \
         2>/dev/null | grep -Poim 1 '^(\h*#)?\h*'"$l_option"'\h*=\h*\H+\b')"
         l_option_value="$(cut -d= -f2 <<< "${l_opt//# /}" | xargs)"
         [ -n "$l_option_value" ] && \
         a_output+=(" - The default value: \"${l_opt//#/}\"" \
         "    is being used in the configuration")
      fi
   fi
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Example output:

```
 - "ProcessSizeMax" is set to: "0" in: "/etc/systemd/coredump.conf.d/60-
coredump.conf"
```

Note:

- **- IF -** the `systemd-coredump` package is not installed this is considered a passing state and remediation is not necessary.
- If multiple lines are returned:
    - The first line includes the value being used by `systemd-coredump`. If this is a correct value, this is considered a passing state.
    - Files that include an incorrect value should be modified to correct or comment out the incorrect value to minimize the potential of the incorrect value being used by `systemd-coredump` due to system configuration changes
    - If the default value is displayed, and is the correct value, this is considered a passing state. However, it is recommended to follow the Remediation procedure

Internal Only - General

to mitigate the potential of a system update changing the default value to an incorrect value.

Remediation:

1. Review all files ending in `.conf` in the `/etc/systemd/coredump.conf.d/` directory and comment out or remove all `ProcessSizeMax=` lines in the `[Coredump]` block that are not `ProcessSizeMax=0`.

Example script:

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

{
   l_block="Coredump" l_option="ProcessSizeMax" l_value="0"
   while IFS= read -r -d $'\0' l_file; do
      awk '/\['"$l_block"'\]/{a=1;next}/\[/{a=0}a' "$l_file" 2>/dev/null | \
      grep -Poi '\h*'"$l_option"'\h*=\h*\H+\b' | \
      grep -Pivq '^\h*'"$l_option"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*[Pp]rocess[Ss]ize[Mm]ax\s*=/s/^/# /' "$l_file"
   done < <(find /etc/systemd/coredump.conf.d/ -type f -name '*.conf' -
print0)
}
```

2. Create or edit a file in the `/etc/systemd/coredump.conf.d` directory ending in `.conf` and edit or add the following line to the `[Coredump]` section:

```
ProcessSizeMax=0
```

Example:

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

{
   [ ! -d /etc/systemd/coredump.conf.d/ ] && mkdir -p
/etc/systemd/coredump.conf.d/
   l_file="/etc/systemd/coredump.conf.d/60-coredump.conf"
   if grep -Psq -- '^\h*\[Coredump\]' "$l_file"; then
      sed -ri '/^\s*\[Coredump\]/a ProcessSizeMax=0' "$l_file"
   else
      printf '%s\n' "" "[Coredump]" "ProcessSizeMax=0" >> "$l_file"
   fi
}
```

3. Run the following command to load the updated configuration:

```
# systemctl reload-or-restart systemd-coredump.socket
```

Default Value:

ProcessSizeMax=2G

Internal Only - General

References:

1. GRID: NIX-00000109
2. CCI-000366
3. NIST SP 800-53 :: CM-6 b
4. NIST SP 800-53A :: CM-6.1 (iv)
5. STIG ID: RHEL-08-010675 | RULE ID: SV-230315r1017126 | CAT II
6. STIG ID: OL08-00-010675 | RULE ID: SV-248633r991589 | CAT II
7. STIG ID: RHEL-09-213085 | RULE ID: SV-257812r1051005 | CAT II
8. STIG ID: ALMA-09-011570 | RULE ID: SV-269181r1050063 | CAT II
9. STIG ID: OL09-00-002381 | RULE ID: SV-271729r1091899 | CAT II
10. COREDUMP.CONF(5)

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

1.5.12 Ensure systemd-coredump Storage is configured (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.

Rationale:

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

Impact:

The core dumps may be logged, but not stored permanently.

Audit:

Run the following script to verify `systemd-coredump Storage` is set to `none`:

Internal Only - General

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

{
   l_analyze_cmd="$(readlink -e /bin/systemd-analyze || \
   readlink -e /usr/bin/systemd-analyze)"
   l_conf_file="systemd/coredump.conf" l_block="Coredump"
   l_option="Storage" l_option_value="" a_output=()
   while IFS= read -r l_file; do
      l_file="${l_file//# /}"
      l_opt="$(awk '/\['"$l_block"'\]/{a=1;next}/\[/{a=0}a' "$l_file" \
      2>/dev/null | grep -Poi '^\h*'"$l_option"'\h*=\h*\H+\b' | tail -n 1)"
         l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_option\" is set to: \"$l_option_value\"" \
      "    in: \"$l_file\"")
   done < <("$l_analyze_cmd" cat-config "$l_conf_file" | tac | \
   grep -Pio '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   if [ "${#a_output[@]}" -le "0" ]; then
      l_file="$(readlink -e /etc/"$l_conf_file" || \
          readlink -e /usr/lib/"$l_conf_file")"
      if [ -e "$l_file" ]; then
         l_opt="$(awk '/\['"$l_block"'\]/{a=1;next}/\[/{a=0}a' "$l_file" \
         2>/dev/null | grep -Poim 1 '^(\h*#)?\h*'"$l_option"'\h*=\h*\H+\b')"
         l_option_value="$(cut -d= -f2 <<< "${l_opt//# /}" | xargs)"
         [ -n "$l_option_value" ] && \
         a_output+=(" - The default value: \"${l_opt//#/}\"" \
         "    is being used in the configuration")
      fi
   fi
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Example output:

```
 - "Storage" is set to: "none" in: "/etc/systemd/coredump.conf.d/60-
coredump.conf"
```

Note:

- **- IF -** the `systemd-coredump` package is not installed this is considered a passing state and remediation is not necessary.
- If multiple lines are returned:
    - The first line includes the value being used by `systemd-coredump`. If this is a correct value, this is considered a passing state.
    - Files that include an incorrect value should be modified to correct or comment out the incorrect value to minimize the potential of the incorrect value being used by `systemd-coredump` due to system configuration changes
    - If the default value is displayed, and is the correct value, this is considered a passing state. However, it is recommended to follow the Remediation procedure

Internal Only - General

to mitigate the potential of a system update changing the default value to an incorrect value.

Remediation:

1. Review all files ending in `.conf` in the `/etc/systemd/coredump.conf.d/` directory and comment out or remove all `Storage=` lines in the `[Coredump]` block that are not `Storage=none`.

Example script:

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

{
   l_block="Coredump" l_option="Storage" l_value="none"
   while IFS= read -r -d $'\0' l_file; do
      awk '/\['"$l_block"'\]/{a=1;next}/\[/{a=0}a' "$l_file" 2>/dev/null | \
      grep -Poi '\h*'"$l_option"'\h*=\h*\H+\b' | \
      grep -Pivq '^\h*'"$l_option"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*[Ss]torage\s*=/s/^/# /' "$l_file"
   done < <(find /etc/systemd/coredump.conf.d/ -type f -name '*.conf' -
print0)
}
```

2. Create or edit a file in the `/etc/systemd/coredump.conf.d` directory ending in `.conf` and edit or add the following line to the `[Coredump]` section:

```
Storage=none
```

Example:

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

{
   [ ! -d /etc/systemd/coredump.conf.d/ ] && mkdir -p \
   /etc/systemd/coredump.conf.d/
   l_file="/etc/systemd/coredump.conf.d/60-coredump.conf"
   if grep -Psq -- '^\h*\[Coredump\]' "$l_file"; then
      sed -ri '/^\s*\[Coredump\]/a Storage=none' "$l_file"
   else
      printf '%s\n' "" "[Coredump]" "Storage=none" >> "$l_file"
   fi
}
```

3. Run the following command to load the updated configuration:

```
# systemctl reload-or-restart systemd-coredump.socket
```

Default Value:

Storage=external

Internal Only - General

References:

1. GRID: NIX-00000110
2. CCI-000366
3. NIST SP 800-53 :: CM-6 b
4. NIST SP 800-53A :: CM-6.1 (iv)
5. STIG ID: RHEL-08-010674 | RULE ID: SV-230314r1017125 | CAT II
6. STIG ID: OL08-00-010674 | RULE ID: SV-248632r991589 | CAT II
7. STIG ID: RHEL-09-213090 | RULE ID: SV-257813r991589 | CAT II
8. STIG ID: ALMA-09-011460 | RULE ID: SV-269180r1050062 | CAT II
9. STIG ID: OL09-00-002382 | RULE ID: SV-271730r1091902 | CAT II
10. COREDUMP.CONF(5)

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

1.6 Configure Command Line Warning Banners

Presenting a warning message prior to the normal user login may assist in the prosecution of trespassers on the computer system. Changing some of these login banners also has the side effect of hiding OS version information and other detailed system information from attackers attempting to target specific exploits at a system.

Guidelines published by the US Department of Defense require that warning messages include at least the name of the organization that owns the system, the fact that the system is subject to monitoring and that such monitoring is in compliance with local statutes, and that use of the system implies consent to such monitoring. It is important that the organization's legal counsel review the content of all messages before any system modifications are made, as these warning messages are inherently site-specific. More information (including citations of relevant case law) can be found at http://www.justice.gov/criminal/cybercrime/

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

Note: The text provided in the remediation actions for these items is intended as an example only. Please edit to include the specific text for your organization as approved by your legal department.

Internal Only - General

1.6.1 Ensure /etc/motd is 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.

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 script to verify `MOTD` files do not contain system information:

```
# grep -Psi -- "(\\\v|\\\r|\\\m|\\\s|\b$(grep '^ID=' /etc/os-release | cut -
d= -f2 | sed -e 's/"//g')\b)" /etc/motd /etc/motd.d/* 2>/dev/null
```

Review the contents of any remaining message-of-the-day files (`/etc/motd` and `/etc/motd.d/*`) and verify they follow local site policy. If a message of the day is not in use, these files need not exist.

Remediation:

Edit the file found in `/etc/motd.d/*` with the appropriate contents according to your site policy, remove any instances of `\m` , `\r` , `\s` , `\v` or references to the `OS platform`

Run the following command to identify files that still contain system information:

```
# grep -Psi -- "(\\\v|\\\r|\\\m|\\\s|\b$(grep '^ID=' /etc/os-release | cut -
d= -f2 | sed -e 's/"//g')\b)" /etc/motd /etc/motd.d/* 2>/dev/null
```

- OR -

Internal Only - General

- IF - the `motd` is not used, this file can be removed.

Run the following command to remove the `motd` file:

```
# rm /etc/motd
```

References:

1. GRID: NIX-00000111
2. NIST SP 800-53 Rev. 5: CM-6, CM-1, CM-3

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

1.6.2 Ensure /etc/issue is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The file `/etc/issue` is read from agetty(8) and similar tools and shown above the login prompt. While some of the information can be adjusted with variables, the file itself is static. To display dynamic content, issue-generator(1) generates a /run/issue file from different files.

Files in `/etc/issue.d` override files with the same name in `/usr/lib/issue.d` and `/run/issue.d`. Files in `/run/issue.d` override files with the same name in /usr/lib/issue.d. Packages should install their configuration files in `/usr/lib/issue.d`. Files in `/etc/issue.d` are reserved for the local administrator, who may use this logic to override the files installed by vendor packages. All configuration files are sorted by their filename in lexicographic order, regardless of which of the directories they reside in.

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 no results are returned:

```
# grep -Psi "(\\\v|\\\r|\\\m|\\\s|\b$(grep '^ID=' /etc/os-release | cut -d= -
f2 | sed -e 's/"//g')\b)" /etc/issue /usr/lib/issue.d/* /etc/issue.d/*
/run/issue.d/*
```

Nothing should be returned

Remediation:

Edit the `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`

Example:

Internal Only - General

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

Note:

- Files in `/etc/issue.d` override files with the same name in `/usr/lib/issue.d` and `/run/issue.d`. Files in `/run/issue.d` override files with the same name in `/usr/lib/issue.d`.
    - Files in `/etc/issue.d` are reserved for the local administrator, who may use this logic to override the files installed by vendor packages.
    - All configuration files are sorted by their filename in lexicographic order, regardless of which of the directories they reside in.

References:

1. GRID: NIX-00000112
2. NIST SP 800-53 Rev. 5: CM-6, CM-1, CM-3

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

1.6.3 Ensure /etc/issue.net is 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.

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 -Psi "(\\v|\\r|\\m|\\s|\b$(grep '^ID=' /etc/os-release | cut -d= -
f2 | sed -e 's/"//g')\b)" /etc/issue.net
```

Note: `/etc/issue.net` may not exist by default, best practice is to create an `/etc/issue.net` file with the appropriate contents according to your site policy and remove any instances of `\m` , `\r` , `\s` , `\v` or references to the `OS platform`

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`

Example:

Internal Only - General

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

References:

1. GRID: NIX-00000113
2. NIST SP 800-53 Rev. 5: CM-6, CM-1, CM-3

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

1.6.4 Ensure pam_motd is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`pam_motd` is a PAM module that can be used to display arbitrary `MOTD` (message of the day) files after a successful login. By default, `pam_motd` shows files in the following locations:

- `/etc/motd`
- `/run/motd`
- `/usr/lib/motd`
- `/etc/motd.d/`
- `/run/motd.d/`
- `/usr/lib/motd.d/`

If `/etc/motd` does not exist, then `/run/motd` is shown. If `/run/motd` does not exist, then `/usr/lib/motd` is shown.

Similar overriding behavior applies to the directories. Files in `/etc/motd.d/` override files with the same name in `/run/motd.d/` and `/usr/lib/motd.d/`. Files in `/run/motd.d/` override files with the same name in `/usr/lib/motd.d/`.

Files in the directories listed above are displayed in lexicographic order by name. Moreover, the files are filtered by reading them with the credentials of the target user authenticating on the system.

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 script to verify that `motd=` is explicitly set for `pam_motd`:

Internal Only - General

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

for svc in sshd login su gdm-password; do
   if [[ -f "/etc/pam.d/$svc" ]]; then
      file="/etc/pam.d/$svc"
  elif [[ -f "/usr/lib/pam.d/$svc" ]]; then
      file="/usr/lib/pam.d/$svc"
  else
      continue
   fi

    while IFS= read -r line; do
       while IFS= read -r motd_path; do
       motd_path="${motd_path%\"}"; motd_path="${motd_path#\"}"
       motd_path="${motd_path%\'}"; motd_path="${motd_path#\'}"

        [[ -r "$motd_path" ]] || continue

        if grep -Pi -- "(\\\v|\\\r|\\\m|\\\s|\b$(grep ^ID= /etc/os-release |
cut -d= -f2 | sed -e 's/"//g')\b)" "$motd_path"; then
           printf '%s\n%s\nmotd=%s\n\n' "$file" "$line" "$motd_path"
        fi
      done < <(printf '%s\n' "$line" | grep -oP
'\bmotd=\K"[^"]+"|\bmotd=\K'\''[^'\'']+'\''|\bmotd=\K\S+')
   done < <(grep -Pi
'^\h*session\h+(required|optional)\h+pam_motd\.so\b.*\bmotd=' "$file"
2>/dev/null)
done
```

Nothing should be returned

Note: Entries in `/etc/pam.d/` overrides vendor defaults in `/usr/lib/pam.d` directories

Remediation:

Edit the `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`

> **Note:** On systems where `/etc/pam.d/sshd` and `/etc/pam.d/login` configure `pam_motd.so` with `motd=/run/motd.dynamic`, the login banner is regenerated at each login by the scripts in `/etc/update-motd.d/`. Review those scripts (e.g. `10-uname`) and remove or edit any that reintroduce OS/kernel/platform-identifying information, then confirm the regenerated banner (`run-parts /etc/update-motd.d/ | tee /run/motd.dynamic`, or a fresh login) complies with your site policy.

References:

1. GRID: NIX-00000388

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

1.6.5 Ensure sshd warning Banner is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The sshd `Banner` parameter specifies a file whose contents must be sent to the remote user before authentication is permitted.

Note: More information about the openSSH server configuration is available in the "Configure SSH Server" section overview.

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.

Not having a properly configured SSH banner, or leaving it at the default value, can expose an organization to legal and compliance risks by failing to provide users with necessary access policies or warnings about unauthorized access. It may violate regulations requiring user acknowledgment of system usage rules, reduce security awareness, and leave systems vulnerable to malicious or unintended access.

Audit:

Run the following command and verify that the contents or the file being called by the `Banner` argument match site policy:

```
# [ -e "$(sshd -T | awk '$1 == "banner" {print $2}')" ] && cat "$(sshd -T |
awk '$1 == "banner" {print $2}')"
```

Run the following command and verify no results are returned:

```
# grep -Psi -- "(\\\v|\\\r|\\\m|\\\s|\b$(grep '^ID=' /etc/os-release | cut -
d= -f2 | sed -e 's/"//g')\b)" "$(sshd -T | awk '$1 == "banner" {print $2}')"
```

Remediation:

Edit the file being called by the `Banner` argument with the appropriate contents according to your site policy, remove any instances of `\m` , `\r` , `\s` , `\v` or references to the `OS platform`

Example:

Internal Only - General

```
# printf '%s\n' "Authorized users only. All activity may be monitored and
reported." > "$(sshd -T | awk '$1 == "banner" {print $2}')"
```

References:

1. GRID: NIX-00000392
2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

1.6.6 Ensure access to /etc/motd is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The contents of the `MOTD` file are displayed to users after login and function as a message of the day for authenticated users.

Rationale:

- IF - the `MOTD` file does not have the correct access configured, it could be modified by unauthorized users with incorrect or misleading information.

Audit:

Run the following command and verify that if `MOTD` file exists, `Access` is `644` or more restrictive, `Uid` and `Gid` are both `0/root`:

```
# stat -Lc 'File: %n  Access: (%a/%A)  Uid: ( %u/ %U)  Gid: ( %g/ %G)' \
   /etc/motd /run/motd /usr/lib/motd \
   /etc/motd.d/* /run/motd.d/* /usr/lib/motd.d/* \
   /etc/update-motd.d/* 2>/dev/null
```

Verify output includes:

```
Access: (0644/-rw-r--r--)  Uid: ( 0/ root) Gid: ( 0/ root)
  -- OR --
Nothing is returned
```

Remediation:

Run the following commands to set mode, owner, and group on the appropriate `MOTD` file:

```
# chown root:root </path/to/motd/file>
# chmod u-x,go-wx </path/to/motd/file>
```

Example

```
# chown root:root $(readlink -e /etc/motd)
# chmod u-x,go-wx $(readlink -e /etc/motd)
```

- OR -

Run the following command to remove the `/etc/motd` file:

Internal Only - General

```
# rm /etc/motd
```

References:

1. GRID: NIX-00000114
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 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.	●	●	●

Internal Only - General

1.6.7 Ensure access to /etc/issue is 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 access configured, it could be modified by unauthorized users with incorrect or misleading information.

Audit:

Run the following command and verify `Access` is `644` or more restrictive and `Uid` and `Gid` are both `0/root`:

```
# stat -Lc 'Access: (%#a/%A)  Uid: ( %u/ %U) Gid: { %g/ %G)' /etc/issue
/usr/lib/issue.d/* /etc/issue.d/* /run/issue.d/* 2>/dev/null
```

Verify output includes:

```
Access: (0644/-rw-r--r--)  Uid: ( 0/ root) Gid: { 0/ root)
```

Remediation:

Run the following commands to set mode, owner, and group on the appropriate `issue` file:

```
# chown root:root <path/to/issue/file>
# chmod u-x,go-wx <path/to/issue/file>
```

Example

```
# chown root:root $(readlink -e /etc/issue)
# chmod u-x,go-wx $(readlink -e /etc/issue)
```

Default Value:

Access: (0644/-rw-r--r--) Uid: (0/ root) Gid: (0/ root)

References:

1. GRID: NIX-00000115
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

Internal Only - General

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.	●	●	●

Internal Only - General

1.6.8 Ensure access to /etc/issue.net is 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 access configured, it could be modified by unauthorized users with incorrect or misleading information.

Audit:

Run the following command and verify `Access` is `644` or more restrictive and `Uid` and `Gid` are both `0/root`:

```
# stat -Lc 'Access: (%#a/%A)  Uid: ( %u/ %U) Gid: { %g/ %G)' /etc/issue.net
```

Verify output includes:

```
Access: (0644/-rw-r--r--)  Uid: ( 0/ root)   Gid: ( 0/ root)
```

Note: `/etc/issue.net` may not exist by default, best practice is to create an `/etc/issue.net` file with the appropriate contents according to your site policy and remove any instances of `\m` , `\r` , `\s` , `\v` or references to the `OS platform`

Remediation:

Run the following commands to set mode, owner, and group on `/etc/issue.net`:

```
# chown root:root $(readlink -e /etc/issue.net)
# chmod u-x,go-wx $(readlink -e /etc/issue.net)
```

Default Value:

Access: (0644/-rw-r--r--) Uid: (0/ root) Gid: (0/ root)

References:

1. GRID: NIX-00000116
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

Internal Only - General

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.	●	●	●

Internal Only - General

1.6.9 Ensure access to pam_motd file is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`pam_motd` is a PAM module that can be used to display arbitrary `motd` (message of the day) files after a successful login. By default, `pam_motd` shows files in the following locations:

```
/etc/motd
/run/motd
/usr/lib/motd
/etc/motd.d/
/run/motd.d/
/usr/lib/motd.d/
```

Rationale:

If the `pam_motd` file does not have the correct access configured, it could be modified by unauthorized users with incorrect or misleading information.

Audit:

Run the following command and verify Access is `0644` or more restrictive and `Uid` and `Gid` are both `0/root`:

```
# grep -hPo 'motd=\K\S+' /etc/pam.d/* 2>/dev/null | sort -u | xargs -r stat -
Lc 'Access: (%#a/%A)  Uid: ( %u/ %U) Gid: ( %g/ %G)'
```

Example Output

```
Access: (0644/-rw-r--r--)  Uid: ( 0/ root) Gid: ( 0/ root)
```

Remediation:

Run the following commands to set mode, owner, and group to the `pam_motd` file:

```
# chown root:root </path/to/motd/>
# chmod u-x,go-wx </path/to/motd/>
```

Note: The audit procedure command `grep -hPo 'motd=\K\S+' /etc/pam.d/* | sort | uniq` output will provide the path to the `pam_motd` file.

Example

Internal Only - General

```
# chown root:root /run/motd.dynamic
# chmod u-x,go-wx /run/motd.dynamic
```

References:

1. GRID: NIX-00000394

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.	●	●	●

Internal Only - General

1.6.10 Ensure access to sshd warning banner is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The sshd warning banner must be sent to the remote user before authentication is permitted.

Rationale:

If the sshd warning banner file does not have correct access permissions, it could be modified by unauthorized users to present incorrect or misleading content. An attacker who can write to the banner file could remove legal warning language required for compliance, or inject content designed to mislead users. The file should be owned by `root:root` and not writable by group or others.

Impact:

No operational impact. Tightening permissions on the banner file to `root:root 644` does not affect SSH connectivity or authentication.

Audit:

Run the following command and verify the sshd warning banner file access is `644` or more restrictive, `Uid` and `Gid` are both `0/root`:

```
# banner=$(sshd -T | grep -Pi -- '^banner\h+\/\H+') && stat -Lc 'Access:
(%#a/%A)  Uid: ( %u/ %U) Gid: { %g/ %G)' $banner 2>/dev/null
```

Verify output includes:

```
Access: (0644/-rw-r--r--)  Uid: ( 0/ root) Gid: ( 0/ root)
  -- OR --
Nothing is returned
```

If no output is returned, the `Banner` directive is not set — verify rule 1.6.5 (Ensure sshd Banner is configured) is also addressed.

Remediation:

Run the following commands to set the correct owner and permissions on the configured banner file. Replace `</path/to/sshd_banner>` with the actual path returned by `sshd -T | grep banner`:

Internal Only - General

```
# chown root:root </path/to/sshd_banner>
# chmod u-x,go-wx </path/to/sshd_banner>
```

Example using the default banner path:

```
# chown root:root /etc/issue.net
# chmod 644 /etc/issue.net
```

References:

1. GRID: NIX-00000393
2. NIST SP 800-53 Rev. 5: CM-6, AC-3

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.	●	●	●

Internal Only - General

1.6.11 Ensure update-notifier-motd.service is not in use (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`update-notifier-motd.service` is a oneshot systemd unit shipped on Ubuntu by the `update-notifier-common` package. It executes `/usr/lib/update-notifier/update-motd-updates-available`, which refreshes the dynamic MOTD fragments in `/var/lib/update-notifier/` (most visibly `updates-available`, the "N packages can be updated, M are security updates" lines surfaced by `pam_motd`).

The same script also drives the `motd-news` integration via `/etc/update-motd.d/50-motd-news`, which contacts Ubuntu's `motd.ubuntu.com` endpoint when `ENABLED=1` in `/etc/default/motd-news`.

Note:

- `update-notifier-motd.service` is not present on Debian 13 — the `update-notifier-common` package is Ubuntu-specific. On systems where the unit is not installed, this recommendation passes by default.
- The service is normally triggered by its companion `update-notifier-motd.timer`; see also the recommendation covering that timer.

Rationale:

Dynamic MOTD content sourced from the package-management subsystem and from upstream "motd-news" feeds discloses host-specific patch state (package counts, security-update counts, distribution name and version) to anyone who can read the MOTD at the login prompt — including unauthenticated SSH probes when sshd is configured to display the MOTD before the banner. This is information disclosure that aids reconnaissance, conflicts with the Section 1.6 goal of presenting a controlled site-policy login banner, and adds an outbound dependency on third-party infrastructure (`motd.ubuntu.com`) that may not be acceptable in regulated or air-gapped environments.

Stopping and masking the service eliminates both the local information-disclosure surface and the upstream-fetch behaviour without removing the package (which may have other administrative uses).

Internal Only - General

Impact:

The "N packages can be updated, M are security updates" lines and any motd-news content will no longer be refreshed by this service. Administrators relying on those lines for at-a-glance patch status should obtain that information from `apt list --upgradable` or a dedicated patch-management tool instead.

Audit:

Run the following command to verify `update-notifier-motd.service` is not enabled:

```
# systemctl is-enabled update-notifier-motd.service 2>/dev/null | grep
'enabled'
```

Nothing should be returned.

Run the following command to verify `update-notifier-motd.service` is not active:

```
# systemctl is-active update-notifier-motd.service 2>/dev/null | grep
'^active'
```

Nothing should be returned.

Note: On systems where `update-notifier-common` is not installed (e.g., Debian 13), neither command will return output and this recommendation is satisfied.

Remediation:

- IF - `update-notifier-motd.service` is installed (Ubuntu), run the following commands to stop and mask the unit:

```
# systemctl stop update-notifier-motd.service
# systemctl mask update-notifier-motd.service
```

Note:

- `mask` is preferred over `disable` because `disable` only removes enablement symlinks; a transient activation (e.g., via the companion timer, or `systemctl start`) will still run the unit. `mask` redirects the unit to `/dev/null` so it cannot be started.
- This recommendation does not require removing the `update-notifier-common` package. If the package is also undesired, it may be removed with `apt purge update-notifier-common`, but that is out of scope for this rule.

Additional Information:

This rule pairs with the recommendation covering `update-notifier-motd.timer`. Both must be masked to fully suppress the dynamic MOTD refresh on Ubuntu — masking only the service leaves the timer firing repeatedly against the masked unit (the timer activation will fail cleanly, but the failure messages clutter the journal); masking only the timer leaves the service available for manual or `pam_motd`-triggered execution.

Companion rules in this subsection:

Internal Only - General

- `1.6.4` (`Ensure pam_motd is configured`) governs the PAM stack that surfaces the dynamic MOTD content this service generates.
- `1.6.6` / `1.6.9` (access to `/etc/motd` and `pam_motd` paths) govern who can modify the dynamic MOTD fragments after they are generated.

This rule, the companion timer rule, and the `motd-news` rule (separately proposed under ticket 23590) together close the dynamic-MOTD surface on Ubuntu. On Debian 13 all three rules pass by default.

CIS Controls:

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

Internal Only - General

1.6.12 Ensure update-notifier-motd.timer is not in use (Automated)

Profile Applicability:

- Level 1 - Server

Description:

`update-notifier-motd.timer` is the systemd timer shipped on Ubuntu by the `update-notifier-common` package. It periodically activates `update-notifier-motd.service`, which in turn refreshes the dynamic MOTD fragments under `/var/lib/update-notifier/` (the "N packages can be updated, M are security updates" content surfaced by `pam_motd`) and drives the `motd-news` integration via `/etc/update-motd.d/50-motd-news`.

Note:

- `update-notifier-motd.timer` is not present on Debian 13 — the `update-notifier-common` package is Ubuntu-specific. On systems where the unit is not installed, this recommendation passes by default.
- The timer is the recurring trigger for the dynamic MOTD refresh. The matching service rule must also be addressed; see also the recommendation covering `update-notifier-motd.service`.

Rationale:

Even when the companion service is masked, an unmasked recurring timer continues to schedule activations, generating journal noise and leaving a re-enablement path open (an attacker or operator with write access to the unit directories can unmask the service and the timer's existing schedule will immediately resume driving it). Masking the timer removes the recurring trigger, ensuring the dynamic MOTD refresh cannot resume without explicit administrative action on both units.

Impact:

The dynamic MOTD content (package update counts, security update counts, motd-news lines) will no longer be refreshed on the timer's schedule. Whatever values were last written to `/var/lib/update-notifier/updates-available` will remain in the MOTD until manually cleared. Sites that want the cached values removed entirely can delete `/var/lib/update-notifier/updates-available` after masking both units; sites that want to retain the legal banner content without the package-count lines should additionally remove or empty `/etc/update-motd.d/90-updates-available` (the MOTD fragment that reads the cached file).

Audit:

Run the following command to verify `update-notifier-motd.timer` is not enabled:

Internal Only - General

```
# systemctl is-enabled update-notifier-motd.timer 2>/dev/null | grep
'enabled'
```

Nothing should be returned.

Run the following command to verify `update-notifier-motd.timer` is not active:

```
# systemctl is-active update-notifier-motd.timer 2>/dev/null | grep '^active'
```

Nothing should be returned.

Note: On systems where `update-notifier-common` is not installed (e.g., Debian 13), neither command will return output and this recommendation is satisfied.

Remediation:

- IF - `update-notifier-motd.timer` is installed (Ubuntu), run the following commands to stop and mask the unit:

```
# systemctl stop update-notifier-motd.timer
# systemctl mask update-notifier-motd.timer
```

Note:

- `mask` is preferred over `disable` because `disable` only removes enablement symlinks; the timer can still be reactivated transiently with `systemctl start update-notifier-motd.timer`. `mask` redirects the unit to `/dev/null` so it cannot be started.
- This recommendation does not require removing the `update-notifier-common` package. If the package is also undesired, it may be removed with `apt purge update-notifier-common`, but that is out of scope for this rule.
- The companion `update-notifier-motd.service` rule must also be addressed. Masking only the timer leaves the service available for manual or `pam_motd`-triggered execution.

Additional Information:

This rule pairs with the recommendation covering `update-notifier-motd.service`. Both must be masked to fully suppress the dynamic MOTD refresh on Ubuntu — masking only the timer leaves the service available for manual or `pam_motd`-triggered execution; masking only the service leaves the timer firing repeatedly against the masked unit (the timer activation will fail cleanly, but the failure messages clutter the journal).

Companion rules in this subsection:

- `1.6.4` (`Ensure pam_motd is configured`) governs the PAM stack that surfaces the dynamic MOTD content this timer's service refreshes.
- `1.6.6` / `1.6.9` (access to `/etc/motd` and `pam_motd` paths) govern who can modify the dynamic MOTD fragments after they are generated.

Internal Only - General

This rule, the companion service rule, and the `motd-news` rule (separately proposed under ticket 23590) together close the dynamic-MOTD surface on Ubuntu. On Debian 13 all three rules pass by default.

CIS Controls:

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

Internal Only - General

1.7 Configure GNOME Display Manager

The GNOME Display Manager (GDM) is a program that manages graphical display servers and handles graphical user logins.

Note: If GDM is not installed on the system, this section can be skipped

Internal Only - General

1.7.1 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.

By using the lockdown mode in dconf, you can prevent users from changing specific settings.

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.

Without locking down the system settings, user settings take precedence over the system settings.

Audit:

1. Run the following commands to verify the login banner configuration is locked:

Run the following command to verify the text banner on the login screen can not be disabled by users:

```
# gsettings writable org.gnome.login-screen banner-message-enable

false
```

Run the following command to verify the text banner on the login screen can not be modified by local users:

```
# gsettings writable org.gnome.login-screen banner-message-text

false
```

Note: If the settings are not locked, the values generated by the `gsettings get` commands below may be user level opposed to system-wide settings. These setting will need to be validated and potentially remediated after they are locked in accordance with the Remediation procedure.

2. Run the following commands to verify the login banner is enabled and meets local site policy.

Internal Only - General

Run the following command to verify that the system-wide text banner on the login screen is enabled:

```
# gsettings get org.gnome.login-screen banner-message-enable

true
```

Run the following command to verify the system-wide text banner on the login screen meets local site policy:

```
# gsettings get org.gnome.login-screen banner-message-text
```

Example output:

```
'Authorized uses only. All activity may be monitored and reported'
```

Remediation:

1. Create or edit the file `/etc/dconf/profile/user` and add the following lines if they do not exist:

```
user-db:user
system-db:local
```

Example:

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

{
   l_dir="/etc/dconf/profile/"
   [ ! -d "$l_dir" ] && mkdir /etc/dconf/profile/
   ! grep -Psq '^\h*user-db:user\b' "$l_dir/user" && \
   printf '%s\n' "" "user-db:user" >> "$l_dir/user"
   ! grep -Psq '^\h*system-db:local\b' "$l_dir/user" && \
   sed -ri '/^\s*user-db:user/a system-db:local' "$l_dir/user"
}
```

2. Run the following command to create the `/etc/dconf/db/local.d/` and `/etc/dconf/db/local.d/locks/` directories if either does not exist:

```
# [ ! -d "/etc/dconf/db/local.d/locks/" ] && mkdir -p
/etc/dconf/db/local.d/locks/
```

3. Create or edit a file in `/etc/dconf/db/local.d/locks/` and add the following lines to lock the login banner configuration:

```
/org/gnome/login-screen/banner-message-enable
/org/gnome/login-screen/banner-message-text
```

Example:

Internal Only - General

```
# printf '%s\n' "" "/org/gnome/login-screen/banner-message-enable" \
"/org/gnome/login-screen/banner-message-text" >>
/etc/dconf/db/local.d/locks/60-banner-message
```

4. Run the following script to comment out any incorrect settings in a local system-wide database keyfile:

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

{
   f_key_file_fix()
   {
      while IFS= read -r -d $'\0' l_file; do
         grep -Psiq -- "^\h*$l_parameter$l_value(\b|\h*$)" "$l_file" && \
         sed -ri '/^\s*'"$l_parameter"'/s/^/# /g' "$l_file"
      done < <(find /etc/dconf/db -mindepth 2 -maxdepth 2 -type f -print0)
   }
   l_parameter="banner-message-enable=" l_value="false"; f_key_file_fix
   l_parameter="banner-message-text=" l_value="(['\"]{2})?"; f_key_file_fix
}
```

5. Create or edit a local keyfile for machine-wide settings in '/etc/dconf/db/local.d/` with the following lines to set the login banner configuration:

```
[org/gnome/login-screen]
banner-message-enable=true
banner-message-text='Type the banner message here.'
```

Example script:

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

{
   l_file="/etc/dconf/db/local.d/60-banner-message"
   l_banner="'Authorized uses only. All activity may be monitored and
reported'"
   a_keyfile=("[org/gnome/login-screen]" "banner-message-enable=true" \
   "banner-message-text=$l_banner")
   if grep -Psq -- '^\h*\[org\/gnome\/login-screen\]' "$l_file"; then
      ! grep -Psiq -- '^\h*banner-message-enable=true\b' "$l_file" && \
      sed -ri '/^\s*\[org\/gnome\/login-screen\]/a banner-message-
enable=true' "$l_file"
      ! grep -Psiq -- '^\h*banner-message-text=[^#\n\r]+' "$l_file" && \
      sed -ri '/^\s*\[org\/gnome\/login-screen\]/a banner-message-
text='"$l_banner"'' "$l_file"
   else
      printf '%s\n' "" "${a_keyfile[@]}" >> "$l_file"
   fi
}
```

6. Run the following command to update the dconf database:

Internal Only - General

```
# dconf update
```

Note:

- banner-message-text should be set in accordance with local site policy
- Users must log out and back in again before the system-wide settings take effect.
- There is no character limit for the banner message. gnome-shell autodetects longer stretches of text and enters two column mode.
- The banner message cannot be read from an external file

Default Value:

disabled

References:

1. GRID: NIX-00000120
2. CCI-000048 | CCI-001384 | CCI-001385 | CCI-001386 | CCI-001387 | CCI-001388
3. NIST SP 800-53 :: AC-8 a | NIST SP 800-53 :: AC-8 c
4. NIST SP 800-53A :: AC-8.1 (ii) | NIST SP 800-53A :: AC-8.2 (i) | NIST SP 800-53A :: AC-8.2 (ii)
5. STIG ID: RHEL-08-010049 | RULE ID: SV-244519r1017326 | CAT II
6. STIG ID: RHEL-08-010050 | RULE ID: SV-230226r1069298 | CAT II
7. STIG ID: RHEL-09-271010 | RULE ID: SV-258012r1014855 | CAT II
8. STIG ID: RHEL-09-171011 | RULE ID: SV-270174r1044831 | CAT II
9. STIG ID: RHEL-09-271015 | RULE ID: SV-258013r1045082 | CAT II
10. STIG ID: ALMA-09-008490 | RULE ID: SV-269156r1050038 | CAT II
11. STIG ID: ALMA-09-008600 | RULE ID: SV-269157r1050039 | CAT II
12. STIG ID: UBTU-22-271010 | RULE ID: SV-260535r958390 | CAT II
13. STIG ID: UBTU-22-271015 | RULE ID: SV-260536r958390 | CAT II
14. STIG ID: UBTU-24-200650 | RULE ID: SV-270692r1066565 | CAT II
15. STIG ID: UBTU-24-200660 | RULE ID: SV-270693r1066568 | CAT II
16. STIG ID: SLES-15-010080 | RULE ID: SV-234808r958586 | CAT II
17. STIG ID: SLES-15-010090 | RULE ID: SV-234809r958586 | CAT II
18. https://help.gnome.org/admin/system-admin-guide/stable/login-banner.html.en

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.

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

1.7.2 Ensure GDM disable-user-list is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

GNOME Display Manager (GDM) handles graphical login for GNOME based systems.

The `disable-user-list` option controls if a list of users is displayed on the login screen

By using the lockdown mode in dconf, you can prevent users from changing specific settings.

Rationale:

Displaying the user list eliminates half of the Userid/Password equation that an unauthorized person would need to log on.

Without locking down the system settings, user settings take precedence over the system settings.

Audit:

1. Run the following command to verify the `disable-user-list` option is locked:

```
# gsettings writable org.gnome.login-screen disable-user-list

false
```

Note: If the setting is not locked, the value generated by the `gsettings get` command below may be user level opposed to system-wide settings. This setting will need to be validated and potentially remediated after it is locked in accordance with the Remediation procedure.

2. Run the following command and to verify that the `disable-user-list` option is enabled:

```
# gsettings get org.gnome.login-screen disable-user-list

true
```

Remediation:

1. Create or edit the file `/etc/dconf/profile/user` and add the following lines if they do not exist:

Internal Only - General

```
user-db:user
system-db:local
```

Example:

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

{
   l_dir="/etc/dconf/profile/"
   [ ! -d "$l_dir" ] && mkdir /etc/dconf/profile/
   ! grep -Psq '^\h*user-db:user\b' "$l_dir/user" && \
   printf '%s\n' "" "user-db:user" >> "$l_dir/user"
   ! grep -Psq '^\h*system-db:local\b' "$l_dir/user" && \
   sed -ri '/^\s*user-db:user/a system-db:local' "$l_dir/user"
}
```

2. Run the following command to create the `/etc/dconf/db/local.d/` and `/etc/dconf/db/local.d/locks/` directories if either does not exist:

```
# [ ! -d "/etc/dconf/db/local.d/locks/" ] && mkdir -p
/etc/dconf/db/local.d/locks/
```

3. Create or edit a file in `/etc/dconf/db/local.d/locks/` and add the following lines:

```
/org/gnome/login-screen/disable-user-list
```

Example:

```
# printf '%s\n' "" "/org/gnome/login-screen/disable-user-list" >>
/etc/dconf/db/local.d/locks/60-login-screen
```

4. Run the following script to comment out any incorrect settings in a local system-wide database keyfile:

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

{
   l_parameter="disable-user-list=" l_value="false"
   while IFS= read -r -d $'\0' l_file; do
      grep -Psiq -- "^\h*$l_parameter$l_value(\b|\h*$)" "$l_file" && \
      sed -ri '/^\s*'"$l_parameter"'/s/^/# /g' "$l_file"
   done < <(find /etc/dconf/db -mindepth 2 -maxdepth 2 -type f -print0)
}
```

5. Create or edit a local keyfile for machine-wide settings in '/etc/dconf/db/local.d/` with the following lines:

```
[org/gnome/login-screen]
disable-user-list=true
```

Example script:

Internal Only - General

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

{
   l_file="/etc/dconf/db/local.d/60-login-screen"
   a_keyfile=("[org/gnome/login-screen]" "disable-user-list=true")
   if grep -Psq -- '^\h*\[org\/gnome\/login-screen\]' "$l_file"; then
      ! grep -Psiq -- '^\h*disable-user-list=true\b' "$l_file" && \
      sed -ri '/^\s*\[org\/gnome\/login-screen\]/a disable-user-list=true'
"$l_file"
      grep -Psiq -- '^\h*disable-user-list=false\b' "$l_file" && \
      sed -ri 's/^\s*(disable-user-list=)(false).*$/\1true/' "$l_file"
   else
      printf '%s\n' "" "${a_keyfile[@]}" >> "$l_file"
   fi
}
```

6. Run the following command to update the dconf databases:

```
# dconf update
```

Note: Users must log out and back in again before the system-wide settings take effect.

Default Value:

false

References:

1. GRID: NIX-00000119
2. CCI-000366
3. NIST SP 800-53 :: CM-6 b
4. NIST SP 800-53A :: CM-6.1 (iv)
5. STIG ID: RHEL-08-020032 | RULE ID: SV-244536r1017343 | CAT II
6. STIG ID: RHEL-09-271115 | RULE ID: SV-258033r1045120 | CAT II
7. STIG ID: ALMA-09-014430 | RULE ID: SV-269207r1050089 | CAT II
8. https://help.gnome.org/admin/system-admin-guide/stable/login-userlist-disable.html.en

Additional Information:

If a different GUI login service is in use and required on the system, consult your documentation to disable displaying the user list

Internal Only - General

CIS Controls:

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

Internal Only - General

1.7.3 Ensure GDM screen lock is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

GNOME Desktop Manager can make the screen lock automatically whenever the user is idle for some amount of time.

- `idle-delay=uint32 {n}` - Number of seconds of inactivity before the screen goes blank
- `lock-delay=uint32 {n}` - Number of seconds after the screen is blank before locking the screen

By using the lockdown mode in dconf, you can prevent users from changing specific settings.

Rationale:

Setting a lock-out value reduces the window of opportunity for unauthorized user access to another user's session that has been left unattended.

Without locking down the system settings, user settings take precedence over the system settings.

Satisfies: SRG-OS-000029-GPOS-00010, SRG-OS-000031-GPOS-00012, SRG-OS-000480-GPOS-00227

Audit:

1. Run the following commands to verify the screen lock settings are locked:

Run the following command to verify `idle-delay` is locked:

```
# gsettings writable org.gnome.desktop.session idle-delay

false
```

Run the following command to verify `lock-delay` is locked:

```
# gsettings writable org.gnome.desktop.screensaver lock-delay

false
```

Note: If the setting is not locked, the value generated by the gsettings get command below may be user level opposed to system-wide settings. This setting will need to be validated and potentially remediated after it is locked in accordance with the Remediation procedure.

Internal Only - General

2. Run the following commands to verify the screen lock settings are configured:

Run the following commands to verify that `lock-delay` is set to `5` or less:

```
# gsettings get org.gnome.desktop.screensaver lock-delay
```

Example output:

```
uint32 5
```

Run the following command to verify `idle-delay` is set to `900` (15 minutes) or less and not `0`:

```
# gsettings get org.gnome.desktop.session idle-delay
```

Example output:

```
uint32 900
```

Remediation:

1. Create or edit the file `/etc/dconf/profile/user` and add the following lines if they do not exist:

```
user-db:user
system-db:local
```

Example:

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

{
   l_dir="/etc/dconf/profile/"
   [ ! -d "$l_dir" ] && mkdir /etc/dconf/profile/
   ! grep -Psq '^\h*user-db:user\b' "$l_dir/user" && \
   printf '%s\n' "" "user-db:user" >> "$l_dir/user"
   ! grep -Psq '^\h*system-db:local\b' "$l_dir/user" && \
   sed -ri '/^\s*user-db:user/a system-db:local' "$l_dir/user"
}
```

2. Run the following command to create the `/etc/dconf/db/local.d/` and `/etc/dconf/db/local.d/locks/` directories if either does not exist:

```
# [ ! -d "/etc/dconf/db/local.d/locks/" ] && mkdir -p
/etc/dconf/db/local.d/locks/
```

3. Create or edit a file in `/etc/dconf/db/local.d/locks/` and add the following lines to lock the login banner configuration:

```
/org/gnome/desktop/session/idle-delay
/org/gnome/desktop/screensaver/lock-delay
```

Example:

Internal Only - General

```
# printf '%s\n' "" "/org/gnome/desktop/session/idle-delay" \
"/org/gnome/desktop/screensaver/lock-delay" >>
/etc/dconf/db/local.d/locks/60-screensaver
```

4. Run the following script to comment out any incorrect settings in a local system-wide database keyfile:

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

{
   f_key_file_fix()
   {
      while IFS= read -r -d $'\0' l_file; do
         grep -Psiq -- "^\h*$l_parameter\h+$l_value(\b|\h*$)" "$l_file" && \
         sed -ri '/^\s*'"$l_parameter"'/s/^/# /g' "$l_file"
      done < <(find /etc/dconf/db -mindepth 2 -maxdepth 2 -type f -print0)
   }
   l_parameter="idle-delay=uint32" l_value="(0|90[1-9]|9[1-9][0-9]|1[0-
9]{3,})"; f_key_file_fix
   l_parameter="lock-delay=uint32" l_value="([6-9]|[1-9][0-9]+)";
f_key_file_fix
}
```

5. Create or edit a local keyfile for machine-wide settings in '/etc/dconf/db/local.d/` with the following lines:

```
[org/gnome/desktop/session]
idle-delay=uint32 900

[org/gnome/desktop/screensaver]
lock-delay=uint32 5
```

- idle-delay=uint32 {n} - Number of seconds of inactivity before the screen goes blank. Should be '900' (15 minutes) or less, and not 0 (disabled)
- lock-delay=uint32 {n} - Number of seconds after the screen is blank before locking the screen. Should be 5 or less

Example script:

Internal Only - General

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

{
   l_file="/etc/dconf/db/local.d/60-screensaver"
   a_keyfile1=("[org/gnome/desktop/session]" "idle-delay=uint32 900")
   a_keyfile2=("[org/gnome/desktop/screensaver]" "lock-delay=uint32 5")
   if grep -Psq -- '^\h*\[org\/gnome\/desktop\/session\]' "$l_file"; then
      ! grep -Psiq -- '^\h*idle-delay=' "$l_file" && \
      sed -ri '/^\s*\[org\/gnome\/desktop\/session\]/a idle-delay=uint32 900'
"$l_file"
      grep -Psiq -- '^\h*idle-delay=uint32\h+(0|90[1-9]|9[1-9][0-9]|[1-9][0-
9]{3,})\b' \
      "$l_file" && sed -ri 's/^\s*(idle-delay=uint32)\s+([0-9]|[1-9][0-
9]+).*$/\1 900/' "$l_file"
   else
      printf '%s\n' "" "${a_keyfile1[@]}" >> "$l_file"
   fi
   if grep -Psq -- '^\h*\[org\/gnome\/desktop\/screensaver\]' "$l_file"; then
      ! grep -Psiq -- '^\h*lock-delay=' "$l_file" && \
      sed -ri '/^\s*\[org\/gnome\/desktop\/screensaver\]/a lock-delay=uint32
5' "$l_file"
      grep -Psiq -- '^\h*lock-delay=uint32\h+([6-9]|[1-9][0-9]+)\b' "$l_file"
&& \
      sed -ri 's/^\s*(lock-delay=uint32)\s+([6-9]|[1-9][0-9]+).*$/\1 5/'
"$l_file"
   else
      printf '%s\n' "" "${a_keyfile2[@]}" >> "$l_file"
   fi
}
```

6. Run the following command to update the dconf database:

```
# dconf update
```

Note: Users must log out and back in again before the system-wide settings take effect.

References:

1. GRID: NIX-00000121
2. CCI-000057
3. NIST SP 800-53 :: AC-11 a
4. NIST SP 800-53A :: AC-11.1 (ii)
5. STIG ID: RHEL-08-020031 | RULE ID: SV-244535r1017342 | CAT II
6. STIG ID: RHEL-08-020060 | RULE ID: SV-230352r1017165 | CAT II
7. STIG ID: RHEL-08-020080 | RULE ID: SV-230354r1069323 | CAT II
8. STIG ID: RHEL-08-020081 | RULE ID: SV-244538r1069324 | CAT II
9. STIG ID: RHEL-09-271065 | RULE ID: SV-258023r958402 | CAT II
10. STIG ID: RHEL-09-271070 | RULE ID: SV-258024r1045100 CAT II
11. STIG ID: RHEL-09-271075 | RULE ID: SV-258025r958402 | CAT II
12. STIG ID: RHEL-09-271080 | RULE ID: SV-258026r1045103 | CAT II
13. STIG ID: ALMA-09-001120 | RULE ID: SV-269103r1049985 | CAT II
14. STIG ID: ALMA-09-001340 | RULE ID: SV-269105r1049987 | CAT II

Internal Only - General

15. STIG ID: ALMA-09-001450 | RULE ID: SV-269106r1049988 | CAT II
16. STIG ID: ALMA-09-001560 | RULE ID: SV-269107r1049989 | CAT II
17. STIG ID: UBTU-22-271025 | RULE ID: SV-260538r958402 | CAT II
18. STIG ID: UBTU-24-200020 | RULE ID: SV-270678r1066523 | CAT II
19. STIG ID: SLES-15-010120 | RULE ID: SV-234812r958402 | CAT II
20. https://help.gnome.org/admin/system-admin-guide/stable/desktop-lockscreen.html.en

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.3 Configure Automatic Session Locking on Enterprise Assets Configure automatic session locking on enterprise assets after a defined period of inactivity. For general purpose operating systems, the period must not exceed 15 minutes. For mobile end-user devices, the period must not exceed 2 minutes.	●	●	●
v7	16.11 Lock Workstation Sessions After Inactivity Automatically lock workstation sessions after a standard period of inactivity.	●	●	●

Internal Only - General

1.7.4 Ensure GDM automount is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 2 - Workstation

Description:

By default GNOME automatically mounts removable media when inserted as a convenience to the user.

By using the lockdown mode in dconf, you can prevent users from changing specific settings.

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.

Without locking down the system settings, user settings take precedence over the system settings.

Satisfies: SRG-OS-000114-GPOS-00059, SRG-OS-000378-GPOS-00163, SRG-OS-000480-GPOS-00227

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:

1. Run the following commands to verify automount settings are locked:

Run the following command to verify `automount` is locked:

```
# gsettings writable org.gnome.desktop.media-handling automount

false
```

Run the following command to verify `automount-open` is locked:

```
# gsettings writable org.gnome.desktop.media-handling automount-open

false
```

2. Run the following commands to verify the automount settings are disabled:

Internal Only - General

Run the following command to verify `automount` is disabled:

```
# gsettings get org.gnome.desktop.media-handling automount

false
```

Run the following command to verify `automount-open` is disabled:

```
# gsettings get org.gnome.desktop.media-handling automount-open

false
```

Remediation:

1. Create or edit the file `/etc/dconf/profile/user` and add the following lines if they do not exist:

```
user-db:user
system-db:local
```

Example:

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

{
   l_dir="/etc/dconf/profile/"
   [ ! -d "$l_dir" ] && mkdir /etc/dconf/profile/
   ! grep -Psq '^\h*user-db:user\b' "$l_dir/user" && \
   printf '%s\n' "" "user-db:user" >> "$l_dir/user"
   ! grep -Psq '^\h*system-db:local\b' "$l_dir/user" && \
   sed -ri '/^\s*user-db:user/a system-db:local' "$l_dir/user"
}
```

2. Run the following command to create the `/etc/dconf/db/local.d/` and `/etc/dconf/db/local.d/locks/` directories if either does not exist:

```
# [ ! -d "/etc/dconf/db/local.d/locks/" ] && mkdir -p
/etc/dconf/db/local.d/locks/
```

3. Create or edit a file in `/etc/dconf/db/local.d/locks/` and add the following lines:

```
/org/gnome/desktop/media-handling/automount
/org/gnome/desktop/media-handling/automount-open
```

Example:

```
# printf '%s\n' "" "/org/gnome/desktop/media-handling/automount" \
"/org/gnome/desktop/media-handling/automount-open" >> \
/etc/dconf/db/local.d/locks/60-media-automount
```

4. Run the following script to comment out any incorrect settings in a local system-wide database keyfile:

Internal Only - General

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

{
   f_key_file_fix()
   {
      while IFS= read -r -d $'\0' l_file; do
         grep -Psiq -- "^\h*$l_parameter$l_value(\b|\h*$)" "$l_file" && \
         sed -ri '/^\s*'"$l_parameter"'/s/^/# /g' "$l_file"
      done < <(find /etc/dconf/db -mindepth 2 -maxdepth 2 -type f -print0)
   }
   l_parameter="automount=" l_value="true"; f_key_file_fix
   l_parameter="automount-open=" l_value="true"; f_key_file_fix
}
```

5. Create or edit a local keyfile for machine-wide settings in '/etc/dconf/db/local.d/` with the following lines:

```
[org/gnome/desktop/media-handling]
automount=false
automount-open=false
```

Example script:

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

{
   l_file="/etc/dconf/db/local.d/60-media-automount"
   a_keyfile=("[org/gnome/desktop/media-handling]" "automount=false"
"automount-open=false")
   if grep -Psq -- '^\h*\[org\/gnome\/desktop\/media-handling\]' "$l_file";
then
      ! grep -Psiq -- '^\h*automount-open=false\b' "$l_file" && \
      sed -ri '/^\s*\[org\/gnome\/desktop\/media-handling\]/a automount-
open=false' "$l_file"
      ! grep -Psiq -- '^\h*automount=false\b' "$l_file" && \
      sed -ri '/^\s*\[org\/gnome\/desktop\/media-handling\]/a
automount=false' "$l_file"
      grep -Psiq -- '^\h*automount=true\b' "$l_file" && \
      sed -ri 's/^\s*(automount=)(true).*$/\1false/' "$l_file"
      grep -Psiq -- '^\h*automount-open=true\b' "$l_file" && \
      sed -ri 's/^\s*(automount-open=)(true).*$/\1false/' "$l_file"
   else
      printf '%s\n' "" "${a_keyfile[@]}" >> "$l_file"
   fi
}
```

6. Run the following command to update the dconf databases:

```
# dconf update
```

Note: Users must log out and back in again before the system-wide settings take effect.

References:

1. GRID: NIX-00000117

Internal Only - General

2. CCI-000778 | CCI-001958
3. NIST SP 800-53 :: IA-3
4. NIST SP 800-53A :: IA-3.1 (ii)
5. STIG ID: RHEL-09-271020 | RULE ID: SV-258014r1045084 | CAT II
6. STIG ID: RHEL-09-271025 | RULE ID: SV-258015r1045086 | CAT II
7. STIG ID: ALMA-09-034890 | RULE ID: SV-269377r1050260 | CAT II
8. STIG ID: ALMA-09-035000 | RULE ID: SV-269378r1050261 | CAT II
9. STIG ID: ALMA-09-035110 | RULE ID: SV-269379r1050262 | CAT II
10. https://access.redhat.com/solutions/20107

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	10.3 Disable Autorun and Autoplay for Removable Media Disable autorun and autoplay auto-execute functionality for removable media.	●	●	●
v7	8.5 Configure Devices Not To Auto-run Content Configure devices to not auto-run content from removable media.	●	●	●

Internal Only - General

1.7.5 Ensure GDM autorun-never is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `autorun-never` setting allows the GNOME Desktop Display Manager to disable autorun through GDM.

By using the lockdown mode in dconf, you can prevent users from changing specific settings.

Rationale:

Malware on removable media may take advantage of Autorun features when the media is inserted into a system and execute.

Without locking down the system settings, user settings take precedence over the system settings.

Satisfies: SRG-OS-000114-GPOS-00059, SRG-OS-000378-GPOS-00163, SRG-OS-000480-GPOS-00227

Audit:

1. Run the following command to verify the `autorun-never` option is locked:

```
# gsettings writable org.gnome.desktop.media-handling autorun-never

false
```

Note: If the setting is not locked, the value generated by the `gsettings get` command below may be user level opposed to system-wide settings. This setting will need to be validated and potentially remediated after it is locked in accordance with the Remediation procedure.

2. Run the following command and to verify that the `autorun-never` option is enabled:

```
# gsettings get org.gnome.desktop.media-handling autorun-never

true
```

Remediation:

1. Create or edit the file `/etc/dconf/profile/user` and add the following lines if they do not exist:

Internal Only - General

```
user-db:user
system-db:local
```

Example:

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

{
   l_dir="/etc/dconf/profile/"
   [ ! -d "$l_dir" ] && mkdir /etc/dconf/profile/
   ! grep -Psq '^\h*user-db:user\b' "$l_dir/user" && \
   printf '%s\n' "" "user-db:user" >> "$l_dir/user"
   ! grep -Psq '^\h*system-db:local\b' "$l_dir/user" && \
   sed -ri '/^\s*user-db:user/a system-db:local' "$l_dir/user"
}
```

2. Run the following command to create the `/etc/dconf/db/local.d/` and `/etc/dconf/db/local.d/locks/` directories if either does not exist:

```
# [ ! -d "/etc/dconf/db/local.d/locks/" ] && mkdir -p
/etc/dconf/db/local.d/locks/
```

3. Create or edit a file in `/etc/dconf/db/local.d/locks/` and add the following lines:

```
/org/gnome/desktop/media-handling/autorun-never
```

Example:

```
# printf '%s\n' "" "/org/gnome/desktop/media-handling/autorun-never" >> \
/etc/dconf/db/local.d/locks/60-media-autorun
```

4. Run the following script to comment out any incorrect settings in a local system-wide database keyfile:

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

{
   l_parameter="autorun-never=" l_value="false"
   while IFS= read -r -d $'\0' l_file; do
      grep -Psiq -- "^\h*$l_parameter$l_value(\b|\h*$)" "$l_file" && \
      sed -ri '/^\s*'"$l_parameter"'/s/^/# /g' "$l_file"
   done < <(find /etc/dconf/db -mindepth 2 -maxdepth 2 -type f -print0)
}
```

5. Create or edit a local keyfile for machine-wide settings in '/etc/dconf/db/local.d/` with the following lines:

```
[org/gnome/desktop/media-handling]
autorun-never=true
```

Example script:

Internal Only - General

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

{
   l_file="/etc/dconf/db/local.d/60-media-autorun"
   a_keyfile=("[org/gnome/desktop/media-handling]" "autorun-never=true")
   if grep -Psq -- '^\h*\[org\/gnome\/desktop\/media-handling\]' "$l_file";
then
      ! grep -Psiq -- '^\h*autorun-never=true\b' "$l_file" && \
      sed -ri '/^\s*\[org\/gnome\/desktop\/media-handling\]/a autorun-
never=true' "$l_file"
      grep -Psiq -- '^\h*autorun-never=false\b' "$l_file" && \
      sed -ri 's/^\s*(autorun-never=)(false).*$/\1true/' "$l_file"
   else
      printf '%s\n' "" "${a_keyfile[@]}" >> "$l_file"
   fi
}
```

6. Run the following command to update the dconf databases:

```
# dconf update
```

Note: Users must log out and back in again before the system-wide settings take effect.

Default Value:

false

References:

1. GRID: NIX-00000118
2. NIST SP 800-53 :: CM-7 (2)
3. STIG ID: RHEL-09-271030 | RULE ID: SV-258016r958804 | CAT II
4. STIG ID: RHEL-09-271035 | RULE ID: SV-258017r1045088 | CAT II
5. CCI-001764

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	10.3 Disable Autorun and Autoplay for Removable Media Disable autorun and autoplay auto-execute functionality for removable media.	●	●	●
v7	8.5 Configure Devices Not To Auto-run Content Configure devices to not auto-run content from removable media.	●	●	●

Internal Only - General

1.7.6 Ensure XDMCP is not enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

X Display Manager Control Protocol (XDMCP) is designed to provide authenticated access to display management services for remote displays

Rationale:

XDMCP is inherently insecure.

- XDMCP is not a ciphered protocol. This may allow an attacker to capture keystrokes entered by a user
- XDMCP is vulnerable to man-in-the-middle attacks. This may allow an attacker to steal the credentials of legitimate users by impersonating the XDMCP server.

Audit:

Run the following script and verify the output:

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

{
   while IFS= read -r l_file; do
      awk '/\[xdmcp\]/{ f = 1;next } /\[/{ f = 0 } f {if
(/^\s*Enable\s*=\s*true/) print "The file: \"'"$l_file"'\" includes: \"" $0
"\" in the \"[xdmcp]\" block"}' "$l_file"
   done < <(grep -Psil -- '^\h*\[xdmcp\]'
/etc/{gdm3,gdm}/{custom,daemon}.conf)
}
```

Nothing should be returned

Remediation:

Edit all files returned by the audit and remove or commend out the `Enable=true` line in the `[xdmcp]` block:

Example File:

Internal Only - General

```
# GDM configuration storage
#
# See /usr/share/gdm/gdm.schemas for a list of available options.

[daemon]
# Uncomment the line below to force the login screen to use Xorg
#WaylandEnable=false

# Enabling automatic login
#  AutomaticLoginEnable = true
#  AutomaticLogin = user1

# Enabling timed login
#  TimedLoginEnable = true
#  TimedLogin = user1
#  TimedLoginDelay = 10

[security]

[xdmcp]
# Enable=true <- **This line should be removed or commented out**

[chooser]

[debug]
# Uncomment the line below to turn on debugging
# More verbose logs
# Additionally lets the X server dump core if it crashes
#Enable=true
```

Default Value:

false

References:

1. GRID: NIX-00000122
2. NIST SP 800-53 Rev. 5: SI-4

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●

Internal Only - General

Controls Version	Control	IG 1	IG 2	IG 3
v7	9.2 Ensure Only Approved Ports, Protocols and Services Are Running Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.		●	●

Internal Only - General

1.7.7 Ensure Xwayland is configured (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Xwayland is a compatibility layer that allows legacy X11 applications to run within a Wayland environment. It's effectively an X server that runs as a Wayland client, enabling existing X11 software to be displayed on a Wayland compositor.

Rationale:

Wayland's security benefits from not relying on X11's network listener. Without X11, there's no network listener, making it harder for malicious actors to exploit vulnerabilities in X11. However, enabling Xwayland (running X11 applications on Wayland) introduces X11's security concerns.

All X vulnerabilities apply to Xwayland, including keylogging, but they only affect X windows and interactions with them.

Malware can potentially exploit Xwayland vulnerabilities to keylog or intercept other input events

Impact:

Many applications haven't been ported to Wayland yet, and Xwayland makes it possible to run these applications without requiring a full switch back to X11. Disabling Xwayland functionality may cause these applications to fail.

Audit:

Run the following commands to verify Xwayland is disabled in the GDM configuration:

```
# for l_file in /etc/gdm3/custom.conf /etc/gdm3/daemon.conf
/etc/gdm/custom.conf /etc/gdm/daemon.conf; do
   [ -e "$l_file" ] && grep -Pi '^\\h*WaylandEnable\\h*=\\h*false\\b'
"$l_file" && echo "WaylandEnable=false found in: $l_file"
done
```

If GDM is not installed this is not applicable.

Verify output includes `WaylandEnable=false` for the relevant configuration file.

Remediation:

Edit the appropriate GDM configuration file and uncomment or add the following line in the `[daemon]` block:

Internal Only - General

- `/etc/gdm3/custom.conf` — if the `gdm3` package is installed
- `/etc/gdm/custom.conf` — if the `gdm` package is installed

```
WaylandEnable=false
```

Example:

```
# GDM configuration storage

[daemon]
WaylandEnable=false

[security]

[xdmcp]

[chooser]

[debug]
# Uncomment the line below to turn on debugging
#Enable=true
```

Default Value:

Enabled

References:

1. GRID: NIX-00000123

CIS Controls:

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

Internal Only - General

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.

Internal Only - General

2.1 Configure Server Services

This section describes services that are installed on systems that specifically need to run these services. If any of these services are not required, it is recommended that the package be removed.

- IF - the package is required for a dependency:

- Ensure the dependent package is approved by local site policy
- Ensure stopping and masking the service and/or socket meets local site policy
- Stop and mask the service and/or socket to reduce the potential attack surface

The following commands can be used to stop and mask the service and socket:

```
# systemctl stop <service_name>.socket <service_name>.service
# systemctl mask <service_name>.socket <service_name>.service
```

Note: This should not be considered a comprehensive list of services not required for normal system operation. You may wish to consider additions to those listed here for your environment

Internal Only - General

2.1.1 Ensure autofs services are not in use (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 the filesystem 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.

There may be packages that are dependent on the `autofs` package. If the `autofs` package is removed, these dependent packages will be removed as well. Before removing the `autofs` package, review any dependent packages to determine if they are required on the system.

- IF - a dependent package is required: stop and mask the `autofs.service` leaving the `autofs` package installed.

Audit:

As a preference `autofs` should not be installed unless other packages depend on it.

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

```
# dpkg-query -s autofs &>/dev/null && echo "autofs is installed"
```

Nothing should be returned.

- OR -

- IF - the package is required for dependencies:

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

Internal Only - General

```
# systemctl is-enabled autofs.service 2>/dev/null | grep 'enabled'

Nothing should be returned
```

Run the following command to verify the `autofs.service` is not active:

```
# systemctl is-active autofs.service 2>/dev/null | grep '^active'

Nothing should be returned
```

Note: If the package is required for a dependency

- Ensure the dependent package is approved by local site policy
- Ensure stopping and masking the service and/or socket meets local site policy

Remediation:

Run the following commands to stop `autofs.service` and remove the `autofs` package:

```
# systemctl stop autofs.service
# apt purge autofs
```

- OR -

- IF - the `autofs` package is required as a dependency:

Run the following commands to stop and mask `autofs.service`:

```
# systemctl stop autofs.service
# systemctl mask autofs.service
```

References:

1. GRID: NIX-00000124
2. NIST SP 800-53 Rev. 5: SI-3, MP-7
3. RHEL 8 STIG Vul ID: V-230502
4. RHEL 8 STIG Rule ID: SV-230502r627750

Additional Information:

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.

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	10.3 Disable Autorun and Autoplay for Removable Media Disable autorun and autoplay auto-execute functionality for removable media.	●	●	●
v7	8.5 Configure Devices Not To Auto-run Content Configure devices to not auto-run content from removable media.	●	●	●

Internal Only - General

2.1.2 Ensure avahi daemon services are not in use (Automated)

Profile Applicability:

- Level 1 - Server
- Level 2 - 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 remove this package to reduce the potential attack surface.

Impact:

There may be packages that are dependent on the `avahi` package. If the `avahi` package is removed, these dependent packages will be removed as well. Before removing the `avahi` package, review any dependent packages to determine if they are required on the system.

- IF - a dependent package is required: stop and mask the `avahi-daemon.socket` and `avahi-daemon.service` leaving the `avahi` package installed.

Audit:

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

```
# dpkg-query -s avahi-daemon &>/dev/null && echo "avahi-daemon is installed"
```

Nothing should be returned.

- OR -

- IF - the `avahi` package is required as a dependency:

Run the following command to verify `avahi-daemon.socket` and `avahi-daemon.service` are not enabled:

Internal Only - General

```
# systemctl is-enabled avahi-daemon.socket avahi-daemon.service 2>/dev/null |
grep 'enabled'

Nothing should be returned
```

Run the following command to verify `avahi-daemon.socket` and `avahi-daemon.service` are not active:

```
# systemctl is-active avahi-daemon.socket avahi-daemon.service 2>/dev/null |
grep '^active'

Nothing should be returned
```

Note: If the package is required for a dependency

- Ensure the dependent package is approved by local site policy
- Ensure stopping and masking the service and/or socket meets local site policy

Remediation:

Run the following commands to stop `avahi-daemon.socket` and `avahi-daemon.service`, and remove the `avahi-daemon` package:

```
# systemctl stop avahi-daemon.socket avahi-daemon.service
# apt purge avahi-daemon
```

- OR -

- IF - the `avahi-daemon` package is required as a dependency:

Run the following commands to stop and mask the `avahi-daemon.socket` and `avahi-daemon.service`:

```
# systemctl stop avahi-daemon.socket avahi-daemon.service
# systemctl mask avahi-daemon.socket avahi-daemon.service
```

References:

1. GRID: NIX-00000125
2. NIST SP 800-53 Rev. 5: SI-4

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●

Internal Only - General

Controls Version	Control	IG 1	IG 2	IG 3
v7	9.2 Ensure Only Approved Ports, Protocols and Services Are Running Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.		●	●

Internal Only - General

2.1.3 Ensure dhcp server services are not in use (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. On Debian 12, the `kea` package provides DHCPv4 and DHCPv6 server services via `kea-dhcp4-server.service`, `kea-dhcp6-server.service`, and `kea-dhcp-ddns-server.service`.

Rationale:

Unless a system is specifically designated as a DHCP server, the `kea` package should be removed to reduce the potential attack surface. A rogue or misconfigured DHCP server can cause network disruption by assigning incorrect IP addresses, DNS servers, or default gateways to clients. If removal is not possible due to a package dependency, all associated services must be stopped and masked to prevent them from running.

Impact:

There may be packages that depend on the `kea` package. Removing it will remove those dependent packages as well. Review any dependent packages before removal to confirm they are not required. **- IF -** a dependent package is required, stop and mask all three `kea` services rather than removing the package.

Audit:

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

```
# dpkg-query -l | awk '{print $2}' | grep -Pio -- 'kea' &>/dev/null && echo
"kea is installed"
```

Nothing should be returned.

- OR -

- IF - the package is required for dependencies:

Run the following command to verify `kea-dhcp-ddns-server.service` `kea-dhcp4-server.service` `kea-dhcp6-server.service` UnitFileState is not enabled and ActiveState is not active:

```
# systemctl show kea-dhcp-ddns-server.service kea-dhcp4-server.service kea-
dhcp6-server.service -p UnitFileState,ActiveState | grep -Pi
'=(enabled|active)'
```

Nothing should be returned.

Internal Only - General

Note: If the package is required for a dependency

- Ensure the dependent package is approved by local site policy
- Ensure stopping and masking the service and/or socket meets local site policy

Remediation:

Run the following commands to stop `kea-dhcp-ddns-server.service` `kea-dhcp4-server.service` `kea-dhcp6-server.service` and remove `kea` package:

```
# systemctl stop kea-dhcp-ddns-server.service kea-dhcp4-server.service kea-dhcp6-server.service
# apt remove kea
```

- OR -

- IF - the `kea` package is required as a dependency:

Run the following commands to stop and mask `dhcpd.service` and `dhcpd6.service`:

```
# systemctl stop kea-dhcp-ddns-server.service kea-dhcp4-server.service kea-dhcp6-server.service
# systemctl mask kea-dhcp-ddns-server.service kea-dhcp4-server.service kea-dhcp6-server.service
```

References:

1. GRID: NIX-00000127
2. dhcpd(8)
3. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

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

Internal Only - General

2.1.4 Ensure dns server services are not in use (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.

Note: `bind9` is the package and `bind.service` is the alias for `named.service`.

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.

Impact:

There may be packages that are dependent on the `bind9` package. If the `bind9` package is removed, these dependent packages will be removed as well. Before removing the `bind9` package, review any dependent packages to determine if they are required on the system.

- IF - a dependent package is required: stop and mask `named.service` leaving the `bind9` package installed.

Audit:

Run the following command to verify `bind9` is not installed:

```
# dpkg-query -s bind9 &>/dev/null && echo "bind9 is installed"
```

Nothing should be returned.

- OR -

- IF - the package is required for dependencies:

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

```
# systemctl is-enabled named.service 2>/dev/null | grep 'enabled'

Nothing should be returned
```

Run the following command to verify the `named.service` is not active:

```
# systemctl is-active named.service 2>/dev/null | grep '^active'

Nothing should be returned
```

Note: If the package is required for a dependency

Internal Only - General

- Ensure the dependent package is approved by local site policy
- Ensure stopping and masking the service and/or socket meets local site policy

Remediation:

Run the following commands to stop `named.service` and remove the `bind9` package:

```
# systemctl stop named.service
# apt purge bind9
```

- OR -

- IF - the `bind9` package is required as a dependency:

Run the following commands to stop and mask `bind9.service`:

```
# systemctl stop named.service
# systemctl mask named.service
```

References:

1. GRID: NIX-00000128
2. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

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

Internal Only - General

2.1.5 Ensure dnsmasq services are not in use (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`dnsmasq` is a lightweight tool that provides DNS caching, DNS forwarding and DHCP (Dynamic Host Configuration Protocol) services.

Rationale:

Unless a system is specifically designated to act as a DNS caching, DNS forwarding and/or DHCP server, it is recommended that the package be removed to reduce the potential attack surface.

Impact:

There may be packages that are dependent on the `dnsmasq` package. If the `dnsmasq` package is removed, these dependent packages will be removed as well. Before removing the `dnsmasq` package, review any dependent packages to determine if they are required on the system.

- IF - a dependent package is required: stop and mask the `dnsmasq.service` leaving the `dnsmasq` package installed.

Audit:

Run one of the following commands to verify `dnsmasq` is not installed:

```
# dpkg-query -s dnsmasq &>/dev/null && echo "dnsmasq is installed"
```

Nothing should be returned.

- OR -

- IF - the package is required for dependencies:

Run the following command to verify `dnsmasq.service` is not enabled:

```
# systemctl is-enabled dnsmasq.service 2>/dev/null | grep 'enabled'

Nothing should be returned
```

Run the following command to verify the `dnsmasq.service` is not active:

```
# systemctl is-active dnsmasq.service 2>/dev/null | grep '^active'

Nothing should be returned
```

Note: If the package is required for a dependency

Internal Only - General

- Ensure the dependent package is approved by local site policy
- Ensure stopping and masking the service and/or socket meets local site policy

Remediation:

Run the following commands to stop `dnsmasq.service` and remove `dnsmasq` package:

```
# systemctl stop dnsmasq.service
# apt purge dnsmasq
```

- OR -

- IF - the `dnsmasq` package is required as a dependency:

Run the following commands to stop and mask the `dnsmasq.service`:

```
# systemctl stop dnsmasq.service
# systemctl mask dnsmasq.service
```

References:

1. GRID: NIX-00000129
2. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

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

Internal Only - General

2.1.6 Ensure ftp server services are not in use (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The File Transfer Protocol (FTP) provides networked computers with the ability to transfer files. `vsftpd` is the Very Secure File Transfer Protocol Daemon.

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.

Impact:

There may be packages that are dependent on the `vsftpd` package. If the `vsftpd` package is removed, these dependent packages will be removed as well. Before removing the `vsftpd` package, review any dependent packages to determine if they are required on the system.

- IF - a dependent package is required: stop and mask the `vsftpd.service` leaving the `vsftpd` package installed.

Audit:

Run the following command to verify the `vsftpd` package is not installed:

```
# dpkg-query -s vsftpd &>/dev/null && echo  "vsftpd is installed"
```

Nothing should be returned.

- OR -

- IF - the package is required for dependencies:

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

```
# systemctl is-enabled vsftpd.service 2>/dev/null | grep 'enabled'
```

Nothing should be returned.

Run the following command to verify the `vsftpd` service is not active:

```
# systemctl is-active vsftpd.service 2>/dev/null | grep '^active'
```

Nothing should be returned.

Internal Only - General

Note:

- Other ftp server packages may exist. They should also be audited, if not required and authorized by local site policy
- If the package is required for a dependency:
  - Ensure the dependent package is approved by local site policy
  - Ensure stopping and masking the service and/or socket meets local site policy

Remediation:

Run the following commands to stop `vsftpd.service` and remove the `vsftpd` package:

```
# systemctl stop vsftpd.service
# apt purge vsftpd
```

- OR -

- IF - the `vsftpd` package is required as a dependency:

Run the following commands to stop and mask the `vsftpd.service`:

```
# systemctl stop vsftpd.service
# systemctl mask vsftpd.service
```

Note: Other ftp server packages may exist. If not required and authorized by local site policy, they should also be removed. If the package is required for a dependency, the service should be stopped and masked.

References:

1. GRID: NIX-00000130
2. NIST SP 800-53 Rev. 5: CM-7

Additional Information:

Additional FTP servers also exist and should be audited.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●

Internal Only - General

Controls Version	Control	IG 1	IG 2	IG 3
v7	9.2 Ensure Only Approved Ports, Protocols and Services Are Running Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.		●	●

Internal Only - General

2.1.7 Ensure ldap server services are not in use (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 removed to reduce the potential attack surface.

Impact:

There may be packages that are dependent on the `slapd` package. If the `slapd` package is removed, these dependent packages will be removed as well. Before removing the `slapd` package, review any dependent packages to determine if they are required on the system.

- IF - a dependent package is required: stop and mask the `slapd.service` leaving the `slapd` package installed.

Audit:

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

```
# dpkg-query -s slapd &>/dev/null && echo "slapd is installed"
```

Nothing should be returned.

- OR -

- IF - the package is required for dependencies:

Run the following command to verify `slapd.service` is not enabled:

```
# systemctl is-enabled slapd.service 2>/dev/null | grep 'enabled'

Nothing should be returned
```

Run the following command to verify `slapd.service` is not active:

```
# systemctl is-active slapd.service 2>/dev/null | grep '^active'

Nothing should be returned
```

Note: If the package is required for a dependency

Internal Only - General

- Ensure the dependent package is approved by local site policy
- Ensure stopping and masking the service and/or socket meets local site policy

Remediation:

Run the following commands to stop `slapd.service` and remove the `slapd` package:

```
# systemctl stop slapd.service
# apt purge slapd
```

- OR -

- IF - the `slapd` package is required as a dependency:

Run the following commands to stop and mask `slapd.service`:

```
# systemctl stop slapd.service
# systemctl mask slapd.service
```

References:

1. GRID: NIX-00000132
2. For more detailed documentation on OpenLDAP, go to the project homepage at http://www.openldap.org.
3. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

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

Internal Only - General

2.1.8 Ensure message access server services are not in use (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`dovecot-imapd` and `dovecot-pop3d` are an open source IMAP and POP3 server for Linux based systems.

Rationale:

Unless POP3 and/or IMAP servers are to be provided by this system, it is recommended that the package be removed to reduce the potential attack surface.

Note: Several IMAP/POP3 servers exist and can use other service names. These should also be audited and the packages removed if not required.

Impact:

There may be packages that are dependent on `dovecot-imapd` and/or `dovecot-pop3d` packages. If `dovecot-imapd` and `dovecot-pop3d` packages are removed, these dependent packages will be removed as well. Before removing `dovecot-imapd` and/or `dovecot-pop3d` packages, review any dependent packages to determine if they are required on the system.

- IF - a dependent package is required: stop and mask `dovecot.socket` and `dovecot.service` leaving `dovecot-imapd` and/or `dovecot-pop3d` packages installed.

Audit:

Run the following command to verify `dovecot-imapd` and `dovecot-pop3d` are not installed:

```
# dpkg-query -s dovecot-imapd &>/dev/null && echo "dovecot-imapd is
installed"
```

Nothing should be returned.

```
# dpkg-query -s dovecot-pop3d &>/dev/null && echo "dovecot-pop3d is
installed"
```

Nothing should be returned.

- OR -

- IF - a package is installed **and** is required for dependencies:

Internal Only - General

Run the following commands to verify dovecot.socket and dovecot.service are not enabled:

```
# systemctl is-enabled dovecot.socket dovecot.service 2>/dev/null | grep
'enabled'

Nothing should be returned
```

Run the following command to verify dovecot.socket and dovecot.service are not active:

```
# systemctl is-active dovecot.socket dovecot.service 2>/dev/null | grep
'^active'

Nothing should be returned
```

Note: If the package is required for a dependency

- Ensure the dependent package is approved by local site policy
- Ensure stopping and masking the service and/or socket meets local site policy

Remediation:

Run one of the following commands to remove dovecot-imapd and dovecot-pop3d:

Run the following commands to stop dovecot.socket and dovecot.service, and remove the dovecot-imapd and dovecot-pop3d packages:

```
# systemctl stop dovecot.socket dovecot.service
# apt purge dovecot-imapd dovecot-pop3d
```

- OR -

- IF - a package is installed **and** is required for dependencies:

Run the following commands to stop and mask dovecot.socket and dovecot.service:

```
# systemctl stop dovecot.socket dovecot.service
# systemctl mask dovecot.socket dovecot.service
```

References:

1. GRID: NIX-00000134
2. NIST SP 800-53 Rev. 5: CM-7

Additional Information:

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.

Internal Only - General

CIS Controls:

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

Internal Only - General

2.1.9 Ensure network file system services are not in use (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-kernel-server` package be removed to reduce the remote attack surface.

Impact:

There may be packages that are dependent on the `nfs-kernel-server` package. If the `nfs-kernel-server` package is removed, these dependent packages will be removed as well. Before removing the `nfs-kernel-server` package, review any dependent packages to determine if they are required on the system.

- IF - a dependent package is required: stop and mask the `nfs-server.service` leaving the `nfs-kernel-server` package installed.

Audit:

Run the following command to verify `nfs-kernel-server` is not installed:

```
# dpkg-query -s nfs-kernel-server &>/dev/null && echo "nfs-kernel-server is
installed"
```

Nothing should be returned.

- OR -

- IF - package is required for dependencies:

Run the following command to verify that the `nfs-server.service` is not enabled:

```
# systemctl is-enabled nfs-server.service 2>/dev/null | grep 'enabled'

Nothing should be returned
```

Run the following command to verify the `nfs-server.service` is not active:

Internal Only - General

```
# systemctl is-active nfs-server.service 2>/dev/null | grep '^active'

Nothing should be returned
```

Note: If the package is required for a dependency

- Ensure the dependent package is approved by local site policy
- Ensure stopping and masking the service and/or socket meets local site policy

Remediation:

Run the following command to stop `nfs-server.service` and remove `nfs-kernel-server` package:

```
# systemctl stop nfs-server.service
# apt purge nfs-kernel-server
```

- OR -

- IF - the `nfs-kernel-server` package is required as a dependency:

Run the following commands to stop and mask the `nfs-server.service`:

```
# systemctl stop nfs-server.service
# systemctl mask nfs-server.service
```

References:

1. GRID: NIX-00000135
2. NIST SP 800-53 Rev. 5: CM-6, CM-7

CIS Controls:

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

Internal Only - General

2.1.10 Ensure nis server services are not in use (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. The NIS client (`ypbind`) was used to bind a machine to an NIS server and receive the distributed configuration files.

Rationale:

`ypserv.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 `ypserv.service` be removed and other, more secure services be used

Impact:

There may be packages that are dependent on the `ypserv` package. If the `ypserv` package is removed, these dependent packages will be removed as well. Before removing the `ypserv` package, review any dependent packages to determine if they are required on the system.

- IF - a dependent package is required: stop and mask the `ypserv.service` leaving the `ypserv` package installed.

Audit:

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

```
# dpkg-query -s ypserv &>/dev/null && echo "ypserv is installed"
```

Nothing should be returned.

- OR -

- IF - the package is required for dependencies:

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

```
# systemctl is-enabled ypserv.service 2>/dev/null | grep 'enabled'

Nothing should be returned
```

Run the following command to verify `ypserv.service` is not active:

Internal Only - General

```
# systemctl is-active ypserv.service 2>/dev/null | grep '^active'

Nothing should be returned
```

Note: If the package is required for a dependency

- Ensure the dependent package is approved by local site policy
- Ensure stopping and masking the service and/or socket meets local site policy

Remediation:

Run the following commands to stop `ypserv.service` and remove `ypserv` package:

```
# systemctl stop ypserv.service
# apt purge ypserv
```

- OR -

- IF - the `ypserv` package is required as a dependency:

Run the following commands to stop and mask `ypserv.service`:

```
# systemctl stop ypserv.service
# systemctl mask ypserv.service
```

References:

1. GRID: NIX-00000136
2. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

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

Internal Only - General

2.1.11 Ensure print server services are not in use (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 removed to reduce the potential attack surface.

Impact:

Removing the cups package, or disabling `cups.socket` and/or `cups.service` will prevent printing from the system, a common task for workstation systems.

There may be packages that are dependent on the `cups` package. If the `cups` package is removed, these dependent packages will be removed as well. Before removing the `cups` package, review any dependent packages to determine if they are required on the system.

- IF - a dependent package is required: stop and mask `cups.socket` and `cups.service` leaving the `cups` package installed.

Audit:

Run the following command to verify `cups` is not Installed:

```
# dpkg-query -s cups &>/dev/null && echo "cups is installed"
```

Nothing should be returned.

- OR -

- IF - the `cups` package is required as a dependency:

Run the following command to verify the `cups.socket` and `cups.service` are not enabled:

Internal Only - General

```
# systemctl is-enabled cups.socket cups.service 2>/dev/null | grep 'enabled'

Nothing should be returned
```

Run the following command to verify the `cups.socket` and `cups.service` are not active:

```
# systemctl is-active cups.socket cups.service 2>/dev/null | grep '^active'

Nothing should be returned
```

Note: If the package is required for a dependency

- Ensure the dependent package is approved by local site policy
- Ensure stopping and masking the service and/or socket meets local site policy

Remediation:

Run the following commands to stop `cups.socket` and `cups.service`, and remove the `cups` package:

```
# systemctl stop cups.socket cups.service
# apt purge cups
```

- OR -

- IF - the `cups` package is required as a dependency:

Run the following commands to stop and mask the `cups.socket` and `cups.service`:

```
# systemctl stop cups.socket cups.service
# systemctl mask cups.socket cups.service
```

References:

1. GRID: NIX-00000138
2. http://www.cups.org
3. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●

Internal Only - General

Controls Version	Control	IG 1	IG 2	IG 3
v7	9.2 Ensure Only Approved Ports, Protocols and Services Are Running Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.		●	●

Internal Only - General

2.1.12 Ensure rpcbind services are not in use (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `rpcbind` utility maps RPC services to the ports on which they listen. RPC processes notify `rpcbind` when they start, registering the ports they are listening on and the RPC program numbers they expect to serve. The client system then contacts `rpcbind` on the server with a particular RPC program number. The `rpcbind.service` redirects the client to the proper port number so it can communicate with the requested service.

Portmapper is an RPC service, which always listens on tcp and udp 111, and is used to map other RPC services (such as nfs, nlockmgr, quotad, mountd, etc.) to their corresponding port number on the server. When a remote host makes an RPC call to that server, it first consults with portmap to determine where the RPC server is listening.

Rationale:

A small request (~82 bytes via UDP) sent to the Portmapper generates a large response (7x to 28x amplification), which makes it a suitable tool for DDoS attacks. If `rpcbind` is not required, it is recommended to remove `rpcbind` package to reduce the potential attack surface.

Impact:

Many of the libvirt packages used by Enterprise Linux virtualization, and the `nfs-utils` package used for The Network File System (NFS), are dependent on the `rpcbind` package. If the `rpcbind` package is removed, these dependent packages will be removed as well. Before removing the `rpcbind` package, review any dependent packages to determine if they are required on the system.

- IF - a dependent package is required: stop and mask the `rpcbind.socket` and `rpcbind.service` leaving the `rpcbind` package installed.

Audit:

Run the following command to verify `rpcbind` package is not installed:

```
# dpkg-query -s rpcbind &>/dev/null && echo "rpcbind is installed"
```

Nothing should be returned.

- OR -

- IF - the `rpcbind` package is required as a dependency:

Internal Only - General

Run the following command to verify `rpcbind.socket` and `rpcbind.service` are not enabled:

```
# systemctl is-enabled rpcbind.socket rpcbind.service 2>/dev/null | grep
'enabled'

Nothing should be returned
```

Run the following command to verify `rpcbind.socket` and `rpcbind.service` are not active:

```
# systemctl is-active rpcbind.socket rpcbind.service 2>/dev/null | grep
'^active'

Nothing should be returned
```

Note: If the package is required for a dependency

- Ensure the dependent package is approved by local site policy
- Ensure stopping and masking the service and/or socket meets local site policy

Remediation:

Run the following commands to stop `rpcbind.socket` and `rpcbind.service`, and remove the `rpcbind` package:

```
# systemctl stop rpcbind.socket rpcbind.service
# apt purge rpcbind
```

- OR -

- IF - the `rpcbind` package is required as a dependency:

Run the following commands to stop and mask the `rpcbind.socket` and `rpcbind.service`:

```
# systemctl stop rpcbind.socket rpcbind.service
# systemctl mask rpcbind.socket rpcbind.service
```

References:

1. GRID: NIX-00000139
2. NIST SP 800-53 Rev. 5: CM-6, CM-7

Internal Only - General

CIS Controls:

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

Internal Only - General

2.1.13 Ensure rsync services are not in use (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

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

Rationale:

`rsync.service` presents a security risk as the `rsync` protocol is unencrypted.

The `rsync` package should be removed to reduce the attack area of the system.

Impact:

There may be packages that are dependent on the `rsync` package. If the `rsync` package is removed, these dependent packages will be removed as well. Before removing the `rsync` package, review any dependent packages to determine if they are required on the system.

- IF - a dependent package is required: stop and mask `rsync.service` leaving the `rsync` package installed.

Audit:

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

```
# dpkg-query -s rsync &>/dev/null && echo "rsync is installed"
```

Nothing should be returned.

- OR -

- IF - the `rsync` package is required as a dependency:

Run the following command to verify `rsync.service` is not enabled:

```
# systemctl is-enabled rsync.service 2>/dev/null | grep 'enabled'

Nothing should be returned
```

Run the following command to verify `rsync.service` is not active:

```
# systemctl is-active rsync.service 2>/dev/null | grep '^active'

Nothing should be returned
```

Note: If the package is required for a dependency

Internal Only - General

- Ensure the dependent package is approved by local site policy
- Ensure stopping and masking the service and/or socket meets local site policy

Remediation:

Run the following commands to stop `rsync.service`, and remove the `rsync` package:

```
# systemctl stop rsync.service
# apt purge rsync
```

- OR -

- IF - the `rsync` package is required as a dependency:

Run the following commands to stop and mask `rsync.service`:

```
# systemctl stop rsync.service
# systemctl mask rsync.service
```

References:

1. GRID: NIX-00000140
2. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

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

Internal Only - General

2.1.14 Ensure samba file server services are not in use (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The Samba daemon allows system administrators to configure their Linux systems to share file systems and directories with Windows desktops. Samba will advertise the file systems and directories via the Server Message Block (SMB) protocol. Windows desktop users will be able to mount these directories and file systems as letter drives on their systems.

Rationale:

If there is no need to mount directories and file systems to Windows systems, then this service should be deleted to reduce the potential attack surface.

Impact:

There may be packages that are dependent on the `samba` package. If the `samba` package is removed, these dependent packages will be removed as well. Before removing the `samba` package, review any dependent packages to determine if they are required on the system.

- IF - a dependent package is required: stop and mask the `smbd.service` leaving the `samba` package installed.

Audit:

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

```
# dpkg-query -s samba &>/dev/null && echo "samba is installed"
```

Nothing should be returned.

- OR -

- IF - the package is required for dependencies:

Run the following command to verify `smbd.service` is not enabled:

```
# systemctl is-enabled smbd.service 2>/dev/null | grep 'enabled'

Nothing should be returned
```

Run the following command to verify the `smbd.service` is not active:

Internal Only - General

```
# systemctl is-active smbd.service 2>/dev/null | grep '^active'

Nothing should be returned
```

Remediation:

Run the following commands to stop `smbd.service` and remove `samba` package:

```
# systemctl stop smbd.service
# apt purge samba
```

- OR -

- IF - the `samba` package is required as a dependency:

Run the following commands to stop and mask the `smbd.service`:

```
# systemctl stop smbd.service
# systemctl mask smbd.service
```

References:

1. GRID: NIX-00000141
2. NIST SP 800-53 Rev. 5: CM-6, CM-7

CIS Controls:

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

Internal Only - General

2.1.15 Ensure snmp services are not in use (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Simple Network Management Protocol (SNMP) is a widely used protocol for monitoring the health and welfare of network equipment, computer equipment and devices like UPSs.

Net-SNMP is a suite of applications used to implement SNMPv1 (RFC 1157), SNMPv2 (RFCs 1901-1908), and SNMPv3 (RFCs 3411-3418) using both IPv4 and IPv6.

Support for SNMPv2 classic (a.k.a. "SNMPv2 historic" - RFCs 1441-1452) was dropped with the 4.0 release of the UCD-snmp package.

The Simple Network Management Protocol (SNMP) server is used to listen for SNMP commands from an SNMP management system, execute the commands or collect the information and then send results back to the requesting system.

Rationale:

The SNMP server can communicate using `SNMPv1`, which transmits data in the clear and does not require authentication to execute commands. `SNMPv3` replaces the simple/clear text password sharing used in `SNMPv2` with more securely encoded parameters. If the the SNMP service is not required, the `snmpd` package should be removed to reduce the attack surface of the system.

Note: If SNMP is required:

- The server should be configured for `SNMP v3` only. `User Authentication` and `Message Encryption` should be configured.
- If `SNMP v2` is **absolutely** necessary, modify the community strings' values.

Impact:

There may be packages that are dependent on the `snmpd` package. If the `snmpd` package is removed, these packages will be removed as well.

Before removing the `snmpd` package, review any dependent packages to determine if they are required on the system. If a dependent package is required, stop and mask the `snmpd.service` leaving the `snmpd` package installed.

Audit:

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

Internal Only - General

```
# dpkg-query -s snmpd &>/dev/null && echo "snmpd is installed"
```

Nothing should be returned.

- OR -

- IF - the package is required for dependencies:

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

```
# systemctl is-enabled snmpd.service 2>/dev/null | grep 'enabled'

Nothing should be returned
```

Run the following command to verify the `snmpd.service` is not active:

```
# systemctl is-active snmpd.service 2>/dev/null | grep '^active'

Nothing should be returned
```

Note: If the package is required for a dependency

- Ensure the dependent package is approved by local site policy
- Ensure stopping and masking the service and/or socket meets local site policy

Remediation:

Run the following commands to stop `snmpd.service` and remove the `snmpd` package:

```
# systemctl stop snmpd.service
# apt purge snmpd
```

- OR - If the package is required for dependencies:

Run the following commands to stop and mask the `snmpd.service`:

```
# systemctl stop snmpd.service
# systemctl mask snmpd.service
```

References:

1. GRID: NIX-00000142
2. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●

Internal Only - General

Controls Version	Control	IG 1	IG 2	IG 3
v7	9.2 Ensure Only Approved Ports, Protocols and Services Are Running Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.		●	●

Internal Only - General

2.1.16 Ensure telnet server services are not in use (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `telnetd` package contains the `telnet daemon`, which accepts connections from users from other systems via the telnet protocol.

Rationale:

The telnet protocol is insecure and unencrypted. The use of an unencrypted transmission medium could allow a user with access to sniff network traffic the ability to steal credentials. The ssh package provides an encrypted session and stronger security.

Impact:

There may be packages that are dependent on the `telnetd` package. If the `telnetd` package is removed, these dependent packages will be removed as well. Before removing the `telnetd` package, review any dependent packages to determine if they are required on the system.

- IF - a dependent package is required: stop and mask `inetutils-inetd.service` leaving the `telnetd` package installed.

Audit:

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

```
# dpkg-query -l | awk '{print $2}' | grep -Pio -- '^telnetd|^telnetd-ssl'
&>/dev/null && echo "telnetd is installed"
```

Nothing should be returned.

- OR -

- IF - the `telnetd` package is required as a dependency:

Run the following command to verify `inetutils-inetd.service` is not enabled:

```
# systemctl is-enabled inetutils-inetd.service 2>/dev/null | grep 'enabled'
```

Nothing should be returned.

Run the following command to verify `inetutils-inetd.service` is not active:

```
# systemctl is-active inetutils-inetd.service 2>/dev/null | grep '^active'
```

Nothing should be returned.

Internal Only - General

Note: If the package is required for a dependency

- Ensure the dependent package is approved by local site policy
- Ensure stopping and masking the service and/or socket meets local site policy

Remediation:

Run the following commands to stop `inetutils-inetd.service`, and remove the `telnetd` package:

```
# systemctl stop inetutils-inetd.service
# apt purge telnetd
```

- OR -

- IF - the `telnetd` package is required as a dependency:

Run the following commands to stop and mask inetutils-inetd.service:

```
# systemctl stop inetutils-inetd.service
# systemctl mask inetutils-inetd.service
```

References:

1. GRID: NIX-00000143

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●
v7	2.6 Address unapproved software Ensure that unauthorized software is either removed or the inventory is updated in a timely manner	●	●	●
v7	9.2 Ensure Only Approved Ports, Protocols and Services Are Running Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.		●	●

Internal Only - General

2.1.17 Ensure tftp server services are not in use (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Trivial File Transfer Protocol (TFTP) is a simple protocol for exchanging files between two TCP/IP machines. TFTP servers allow connections from a TFTP Client for sending and receiving files.

Rationale:

Unless there is a need to run the system as a TFTP server, it is recommended that the package be removed to reduce the potential attack surface.

TFTP does not have built-in encryption, access control or authentication. This makes it very easy for an attacker to exploit TFTP to gain access to files

Impact:

TFTP is often used to provide files for network booting such as for PXE based installation of servers.

There may be packages that are dependent on the `tftpd-hpa` package. If the `tftpd-hpa` package is removed, these dependent packages will be removed as well. Before removing the `tftpd-hpa` package, review any dependent packages to determine if they are required on the system.

- IF - a dependent package is required: stop and mask `tftpd-hpa.service` leaving the `tftpd-hpa` package installed.

Audit:

Run the following command to verify `tftpd-hpa` is not installed:

```
# dpkg-query -s tftpd-hpa &>/dev/null && echo "tftpd-hpa is installed"
```

Nothing should be returned.

- OR -

- IF - the package is required for dependencies:

Run the following command to verify `tftpd-hpa.service` is not enabled:

```
# systemctl is-enabled tftpd-hpa.service 2>/dev/null | grep 'enabled'

Nothing should be returned
```

Run the following command to verify the `tftpd-hpa.service` is not active:

Internal Only - General

```
# systemctl is-active tftpd-hpa.service 2>/dev/null | grep '^active'

Nothing should be returned
```

Note: If the package is required for a dependency

- Ensure the dependent package is approved by local site policy
- Ensure stopping and masking the service and/or socket meets local site policy

Remediation:

Run the following commands to stop `tftpd-hpa.service`, and remove the `tftpd-hpa` package:

```
# systemctl stop tftpd-hpa.service
# apt purge tftpd-hpa
```

- OR -

- IF - the `tftpd-hpa` package is required as a dependency:

Run the following commands to stop and mask `tftpd-hpa.service`:

```
# systemctl stop tftpd-hpa.service
# systemctl mask tftpd-hpa.service
```

References:

1. GRID: NIX-00000144
2. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

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

Internal Only - General

2.1.18 Ensure web proxy server services are not in use (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

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

Rationale:

Unless a system is specifically set up to act as a proxy server, it is recommended that the squid package be removed to reduce the potential attack surface.

Note: Several HTTP proxy servers exist. These should be checked and removed unless required.

Impact:

There may be packages that are dependent on the `squid` package. If the `squid` package is removed, these dependent packages will be removed as well. Before removing the `squid` package, review any dependent packages to determine if they are required on the system.

- IF - a dependent package is required: stop and mask the `squid.service` leaving the `squid` package installed.

Audit:

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

```
# dpkg-query -s squid &>/dev/null && echo "squid is installed"
```

Nothing should be returned.

- OR -

- IF - the package is required for dependencies:

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

```
# systemctl is-enabled squid.service 2>/dev/null | grep 'enabled'

Nothing should be returned
```

Run the following command to verify the `squid.service` is not active:

Internal Only - General

```
# systemctl is-active squid.service 2>/dev/null | grep '^active'

Nothing should be returned
```

Note: If the package is required for a dependency

- Ensure the dependent package is approved by local site policy
- Ensure stopping and masking the service and/or socket meets local site policy

Remediation:

Run the following commands to stop `squid.service` and remove the `squid` package:

```
# systemctl stop squid.service
# apt purge squid
```

- OR - If the `squid` package is required as a dependency:

Run the following commands to stop and mask the `squid.service`:

```
# systemctl stop squid.service
# systemctl mask squid.service
```

References:

1. GRID: NIX-00000145
2. NIST SP 800-53 Rev. 5: CM-7

Additional Information:

Several HTTP proxy servers exist. These and other services should be checked

CIS Controls:

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

Internal Only - General

2.1.19 Ensure web server services are not in use (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Web servers provide the ability to host web site content.

Rationale:

Unless there is an approved local site requirement to host web content, web server packages should be removed to reduce the potential attack surface. Web servers introduce listening network services, process untrusted remote input, and commonly expose additional attack surface through modules and plugins. If a web server package cannot be removed due to a package dependency, all associated services and sockets must be stopped and masked.

Impact:

Removal of a web server package will prevent the system from hosting web services. If a web server package is required for a dependency, stopping and masking its service and socket prevents it from running while retaining the package. If the masked service is later legitimately required, it must be unmasked before it can be started.

Audit:

Run the following command to verify `apache2` is not installed:

```
# dpkg-query -s apache2 &>/dev/null && echo "apache2 is installed"
```

Nothing should be returned.

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

```
# dpkg-query -s nginx &>/dev/null && echo "nginx is installed"
```

Nothing should be returned.

- OR -

- IF - a package is installed **and** is required for dependencies:

- IF - the apache2 package is required for dependencies:

Run the following command to verify `apache2.socket` and `apache2.service` `UnitFileState` is not `enabled` and `ActiveState` is not `active`:

Internal Only - General

```
# systemctl show apache2.socket apache2.service -p UnitFileState,ActiveState
| grep -Pi '=(enabled|active)'

Nothing should be returned
```

- IF - the nginx package is required for dependencies:

Run the following command to verify `nginx.service` `UnitFileState` is not `enabled` and `ActiveState` is not `active`:

```
# systemctl show nginx.service -p UnitFileState,ActiveState | grep -Pi
'=(enabled|active)'

Nothing should be returned
```

Note: Other web server packages may exist on the system. They should be audited in the same manner if not required and authorized by local site policy.

Remediation:

Run the following commands to stop `apache2.socket`, `apache2.service` and remove the `apache2` package:

```
# systemctl stop apache2.socket apache2.service
# apt purge apache2
```

Run the following commands to stop `nginx.service` and remove the `nginx` package:

```
# systemctl stop nginx.service
# apt purge nginx
```

- OR -

- IF - a package is installed **and** is required for dependencies:

- IF - the apache2 package is required for dependencies:

Run the following commands to stop and mask `apache2.socket` and `apache2.service`:

```
# systemctl stop apache2.socket apache2.service
# systemctl mask apache2.socket apache2.service
```

- IF - the nginx package is required for dependencies:

Run the following commands to stop and mask `nginx.service`:

```
# systemctl stop nginx.service
# systemctl mask nginx.service
```

Note: Other web server packages may exist. If not required and authorized by local site policy, they should also be removed or masked.

References:

1. GRID: NIX-00000146
2. NIST SP 800-53 :: CM-7

Internal Only - General

CIS Controls:

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

Internal Only - General

2.1.20 Ensure xinetd services are not in use (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 to reduce the attack surface are of the system.

Note: If an `xinetd` service or services are required, ensure that any `xinetd` service not required is stopped and masked

Impact:

There may be packages that are dependent on the `xinetd` package. If the `xinetd` package is removed, these dependent packages will be removed as well. Before removing the `xinetd` package, review any dependent packages to determine if they are required on the system.

-IF- a dependent package is required: stop and mask xinetd.service leaving the xinetd package installed.

Audit:

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

```
# dpkg-query -s xinetd &>/dev/null && echo "xinetd is installed"
```

Nothing should be returned.

-OR-

-IF- the `xinetd` package is required as a dependency:

Run the following command to verify `xinetd.service` is not enabled:

```
# systemctl is-enabled xinetd.service 2>/dev/null | grep 'enabled'

Nothing should be returned
```

Run the following command to verify `xinetd.service` is not active:

Internal Only - General

```
# systemctl is-active xinetd.service 2>/dev/null | grep '^active'

Nothing should be returned
```

Note: If the package is required for a dependency

- Ensure the dependent package is approved by local site policy
- Ensure stopping and masking the service and/or socket meets local site policy

Remediation:

Run the following commands to stop `xinetd.service`, and remove the `xinetd` package:

```
# systemctl stop xinetd.service
# apt purge xinetd
```

-OR-

-IF- the `xinetd` package is required as a dependency:

Run the following commands to stop and mask the `xinetd.service`:

```
# systemctl stop xinetd.service
# systemctl mask xinetd.service
```

References:

1. GRID: NIX-00000148
2. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

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

Internal Only - General

2.1.21 Ensure X window server services are not in use (Automated)

Profile Applicability:

- Level 2 - Server

Description:

The X Window System provides a Graphical User Interface (GUI) where users can have multiple windows in which to run programs and various add on. The X Windows system is typically used on workstations where users login, but not on servers where users typically do not login.

Rationale:

Unless your organization specifically requires graphical login access via X Windows, remove it to reduce the potential attack surface.

Impact:

If a Graphical Desktop Manager (GDM) is in use on the system, there may be a dependency on the `xorg-x11-server-common` package. If the GDM is required and approved by local site policy, the package should **not** be removed.

Many Linux systems run applications which require a Java runtime. Some Linux Java packages have a dependency on specific X Windows xorg-x11-fonts. One workaround to avoid this dependency is to use the "headless" Java packages for your specific Java runtime.

Audit:

- IF - a Graphical Desktop Manager or X-Windows server is not required and approved by local site policy:

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

```
dpkg-query -s xserver-common &>/dev/null && echo "xserver-common is
installed"
```

Nothing should be returned

Remediation:

- IF - a Graphical Desktop Manager or X-Windows server is not required and approved by local site policy:

Run the following command to remove the X Windows Server package:

Internal Only - General

```
# apt purge xserver-common
```

References:

1. GRID: NIX-00000147
2. NIST SP 800-53 Rev. 5: CM-11

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●
v7	2.6 Address unapproved software Ensure that unauthorized software is either removed or the inventory is updated in a timely manner	●	●	●

Internal Only - General

2.1.22 Ensure mail transfer agents are 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.

Impact:

This recommendation is designed around the postfix mail server. Depending on your environment you may have an alternative MTA installed such as exim4 or sendmail. If this is the case consult the documentation for your installed MTA to configure the recommended state.

Audit:

Run the following script to verify that the MTA is not listening on any non-loopback address (`127.0.0.1` or `::1`)

Internal Only - General

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

{
   a_output=(); a_output2=(); a_port_list=("25" "465" "587")
   for l_port_number in "${a_port_list[@]}"; do
      if ss -plntu | grep -P -- ':'"$l_port_number"'\h' | grep -Pvq --
'\h+(127\.0\.0\.1|\[?::1\]?):'"$l_port_number"'\b'; then
         a_output2+=(" - Port \"$l_port_number\" is listening on a non-
loopback network interface")
      else
         a_output+=(" - Port \"$l_port_number\" is not listening on a non-
loopback network interface")
      fi
   done
   if command -v postconf &> /dev/null; then
      l_interfaces="$(postconf -n inet_interfaces)"
   elif command -v exim &> /dev/null; then
      l_interfaces="$(exim -bP local_interfaces)"
   elif command -v sendmail &> /dev/null; then
      l_interfaces="$(grep -i "O DaemonPortOptions=" /etc/mail/sendmail.cf |
grep -oP '(?<=Addr=)[^,+]+' | grep -v '^127\.0\.0\.1$')"
   fi
   if [ -n "$l_interfaces" ]; then
      if grep -Pqi '\ball\b' <<< "$l_interfaces"; then
         a_output2+=(" - MTA is bound to all network interfaces")
      elif ! grep -Pqi '(inet_interfaces\h*=\h*)?(0\.0\.0\.0|::1|loopback-
only)' <<< "$l_interfaces"; then
         a_output2+=(" - MTA is bound to a network interface" "
\"$l_interfaces\"")
      else
         a_output+=(" - MTA is not bound to a non loopback network interface"
"   \"$l_interfaces\"")
      fi
   else
      a_output+=(" - MTA not detected or in use")
   fi
   if [ "${#a_output2[@]}" -le 0 ]; then
      printf '%s\n' "" "- Audit Result:" "  ** PASS **" "${a_output[@]}"
   else
      printf '%s\n' "" "- Audit Result:" "  ** FAIL **" "  * Reasons for
audit failure *" "${a_output2[@]}" ""
      [ "${#a_output[@]}" -gt 0 ] && printf '%s\n' "- Correctly set:"
"${a_output[@]}"
   fi
}
```

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`:

Internal Only - General

```
# systemctl restart postfix
```

Note:

- This recommendation is designed around the postfix mail server.
- Depending on your environment you may have an alternative MTA installed such as exim4 or sendmail. If this is the case consult the documentation for your installed MTA to configure the recommended state.

References:

1. GRID: NIX-00000133
2. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

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

Internal Only - General

2.1.23 Ensure only approved services are listening on a network interface (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

A network port is identified by its number, the associated IP address, and the type of the communication protocol such as TCP or UDP.

A listening port is a network port on which an application or process listens on, acting as a communication endpoint.

Each listening port can be open or closed (filtered) using a firewall. In general terms, an open port is a network port that accepts incoming packets from remote locations.

Rationale:

Services listening on the system pose a potential risk as an attack vector. These services should be reviewed, and if not required, the service should be stopped, and the package containing the service should be removed. If required packages have a dependency, the service should be stopped and masked to reduce the attack surface of the system.

Impact:

There may be packages that are dependent on the service's package. If the service's package is removed, these dependent packages will be removed as well. Before removing the service's package, review any dependent packages to determine if they are required on the system.

- IF - a dependent package is required: stop and mask the `<service_name>.socket` and `<service_name>.service` leaving the service's package installed.

Audit:

Run the following command:

```
# ss -plntu
```

Review the output to ensure:

- All services listed are required on the system and approved by local site policy.
- Both the port and interface the service is listening on are approved by local site policy.
- If a listed service is not required:
    - Remove the package containing the service

Internal Only - General

- **- IF -** the service's package is required for a dependency, stop and mask the service and/or socket

Remediation:

Run the following commands to stop the service and remove the package containing the service:

```
# systemctl stop <service_name>.socket <service_name>.service
# apt purge <package_name>
```

- OR - If required packages have a dependency:

Run the following commands to stop and mask the service and socket:

```
# systemctl stop <service_name>.socket <service_name>.service
# systemctl mask <service_name>.socket <service_name>.service
```

Note: replace `<service_name>` with the appropriate service name.

References:

1. GRID: NIX-00000137
2. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

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

Internal Only - General

2.2 Configure Client Services

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.

Internal Only - General

2.2.1 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 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 `nis` is not installed. Use the following command to provide the needed information:

```
# dpkg-query -s nis &>/dev/null && echo "nis is installed"
```

Nothing should be returned.

Remediation:

Uninstall `nis`:

```
# apt purge nis
```

References:

1. GRID: NIX-00000151
2. NIST SP 800-53 Rev. 5: CM-7, CM-11

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●
v7	2.6 Address unapproved software Ensure that unauthorized software is either removed or the inventory is updated in a timely manner	●	●	●

Internal Only - General

2.2.2 Ensure rsh client is not installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `rsh-client` package contains the client commands for the rsh services.

Rationale:

These legacy clients contain numerous security exposures and have been replaced with the more secure SSH package. Even if the server is removed, it is best to ensure the clients are also removed to prevent users from inadvertently attempting to use these commands and therefore exposing their credentials. Note that removing the `rsh-client` package removes the clients for `rsh` , `rcp` and `rlogin` .

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 `rsh-client` is not installed. Use the following command to provide the needed information:

```
# dpkg-query -s rsh-client &>/dev/null && echo "rsh-client is installed"
```

Nothing should be returned.

Remediation:

Uninstall `rsh`:

```
# apt purge rsh-client
```

References:

1. GRID: NIX-00000269
2. NIST SP 800-53 Rev. 5: CM-7

Internal Only - General

CIS Controls:

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

Internal Only - General

2.2.3 Ensure talk client is not installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `talk` software makes it possible for users to send and receive messages across systems through a terminal session. The `talk` client, which allows initialization of talk sessions, is installed by default.

Rationale:

The software presents a security risk as it uses unencrypted protocols for communication.

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 `talk` is not installed. The following command may provide the needed information:

```
# dpkg-query -s talk &>/dev/null && echo "talk is installed"
```

Nothing should be returned.

Remediation:

Uninstall `talk`:

```
# apt purge talk
```

References:

1. GRID: NIX-00000293
2. NIST SP 800-53 Rev. 5: CM-7

Internal Only - General

CIS Controls:

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

Internal Only - General

2.2.4 Ensure telnet client is not installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `telnet` and `inetutils-telnet` packages contain the `telnet` client, which allows users to start connections to other systems via the telnet protocol.

Rationale:

The `telnet` protocol is insecure and unencrypted. The use of an unencrypted transmission medium could allow an unauthorized user to steal credentials. The `ssh` package provides an encrypted session and stronger security and is included in most Linux distributions.

Impact:

Many insecure service clients are used as troubleshooting tools and in testing environments. Uninstalling them can inhibit capability to test and troubleshoot. If they are required it is advisable to remove the clients after use to prevent accidental or intentional misuse.

Audit:

Run the following command to verify that none of the telnet client packages are installed:

```
# for pkg in telnet inetutils-telnet telnet-ssl; do
 dpkg-query -s "$pkg" 2>/dev/null | grep -q '^Status: install ok installed' \
 && echo "$pkg is installed"
 done
```

Nothing should be returned.

Remediation:

Run the following command to uninstall and purge the telnet client packages:

```
# apt purge telnet inetutils-telnet telnet-ssl
```

References:

1. GRID: NIX-00000152
2. NIST SP 800-53 Rev. 5: CM-7, CM-11

Internal Only - General

CIS Controls:

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

Internal Only - General

2.2.5 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 that `ldap-utils` is not installed. Use the following command to provide the needed information:

```
# dpkg-query -s ldap-utils &>/dev/null && echo "ldap-utils is installed"
```

Nothing should be returned.

Remediation:

Uninstall `ldap-utils`:

```
# apt purge ldap-utils
```

References:

1. GRID: NIX-00000150
2. NIST SP 800-53 Rev. 5: CM-7

Internal Only - General

CIS Controls:

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

Internal Only - General

2.2.6 Ensure ftp client is not installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Several FTP-family client packages may be available on Debian / Ubuntu systems:

- `ftp` / `tnftp` — `tnftp` is an enhanced FTP client; `ftp` is the historic Internet-standard FTP client. Both provide interactive file transfer to remote network sites.
- `inetutils-ftp` — the FTP client from the GNU inetutils suite, an alternative to `tnftp`.
- `ftp-ssl` — a TLS-wrapped FTP client (legacy; superseded by FTPS-capable replacements).
- `atftp` — the Advanced Trivial File Transfer Protocol client. Although TFTP differs from full FTP, the package family historically belongs to the "FTP client" category and is commonly removed alongside the others as part of the same hardening intent.

Unless there is a documented need to run any of these clients (for example, to retrieve build artifacts from a legacy FTP-only mirror), it is recommended that all FTP-family client packages be removed to reduce the potential attack surface and prevent users from inadvertently sending credentials over an unencrypted transport.

Rationale:

Unless there is a need to run the system using Internet standard File Transfer Protocol (for example, to allow anonymous downloads), it is recommended that the package be removed to reduce the potential attack surface.

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. Modern alternatives (`curl`, `wget`, `sftp`, `scp`, `rsync` over SSH) cover the legitimate file-transfer use cases without exposing credentials or data in cleartext.

Audit:

Run the following command to verify that none of the FTP-family client packages are installed:

Internal Only - General

```
# for pkg in ftp tnftp atftp inetutils-ftp ftp-ssl; do
    dpkg-query -s "$pkg" 2>/dev/null | grep -q '^Status: install ok
installed' \
        && echo "$pkg is installed"
  done
```

Nothing should be returned.

Remediation:

Run the following command to uninstall and purge every FTP-family client package that is installed:

```
# apt purge ftp tnftp atftp inetutils-ftp ftp-ssl
```

References:

1. GRID: NIX-00000149
2. NIST SP 800-53 Rev. 5: CM-7, CM-11

CIS Controls:

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

2.3 Configure Time Synchronization

It is recommended that systems be configured to synchronize their time using a service such as `systemd-timesyncd`, or `chrony`.

Virtual systems may be configured to receive their time synchronization from their host system.

The host system must be configured to synchronize its time from an authoritative source to be considered compliant with this section.

Any "physical" clock present on a system should be synchronized from an authoritative time source.

Internal Only - General

Only one time synchronization method should be in use on the system

Notes: Only the section related to the time synchronization method in use on the system should be followed, all other time synchronization recommendations should be skipped

Internal Only - General

2.3.1 Ensure time synchronization is in use

It is recommended that physical systems and virtual guests lacking direct access to the physical host's clock be configured to synchronize their time using a service such as `systemd-timesyncd`, or `chrony`.

Internal Only - General

2.3.1.1 Ensure a single time synchronization daemon is in use (Automated)

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.

Note:

- **On virtual systems where host based time synchronization is available consult your virtualization software documentation and verify that host based synchronization is in use and follows local site policy. In this scenario, this section should be skipped**
- Only **one** time synchronization method should be in use on the system. Configuring multiple time synchronization methods could lead to unexpected or unreliable results

Rationale:

Time synchronization is important to support time sensitive security mechanisms and ensures log files have consistent time records across the enterprise, which aids in forensic investigations.

Audit:

On physical systems, and virtual systems where host based time synchronization is not available.

One of the two time synchronization daemons should be available; `chrony` or `systemd-timesyncd`

Run the following script to verify that a single time synchronization daemon is available on the system:

Internal Only - General

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

{
    l_output="" l_output2=""
    service_not_enabled_chk()
    {
        l_out2=""
        if systemctl is-enabled "$l_service_name" 2>/dev/null | grep -q
'enabled'; then
            l_out2="$l_out2\n  - Daemon: \"$l_service_name\" is enabled on
the system"
        fi
        if systemctl is-active "$l_service_name" 2>/dev/null | grep -q
'^active'; then
            l_out2="$l_out2\n  - Daemon: \"$l_service_name\" is active on the
system"
        fi
    }
    l_service_name="systemd-timesyncd.service" # Check systemd-timesyncd
daemon
    service_not_enabled_chk
    if [ -n "$l_out2" ]; then
        l_timesyncd="y"
        l_out_tsd="$l_out2"
    else
        l_timesyncd="n"
        l_out_tsd="\n  - Daemon: \"$l_service_name\" is not enabled and not
active on the system"
    fi
    l_service_name="chrony.service" # Check chrony
    service_not_enabled_chk
    if [ -n "$l_out2" ]; then
        l_chrony="y"
        l_out_chrony="$l_out2"
    else
        l_chrony="n"
        l_out_chrony="\n  - Daemon: \"$l_service_name\" is not enabled and
not active on the system"
    fi
    l_status="$l_timesyncd$l_chrony"
    case "$l_status" in
        yy)
            l_output2=" - More than one time sync daemon is in use on the
system$l_out_tsd$l_out_chrony"
            ;;
        nn)
            l_output2=" - No time sync daemon is in use on the
system$l_out_tsd$l_out_chrony"
            ;;
        yn|ny)
            l_output=" - Only one time sync daemon is in use on the
system$l_out_tsd$l_out_chrony"
            ;;
        *)
            l_output2=" - Unable to determine time sync daemon(s) status"
            ;;
    esac
```


Internal Only - General

```
    if [ -z "$l_output2" ]; then
        echo -e "\n- Audit Result:\n  ** PASS **\n$l_output\n"
    else
        echo -e "\n- Audit Result:\n  ** FAIL **\n - * Reasons for audit
failure * :\n$l_output2\n"
    fi
}
```

Note: Follow the guidance in the subsection for the time synchronization daemon available on the system and skip the other time synchronization daemon subsection.

Remediation:

On physical systems, and virtual systems where host based time synchronization is not available.

Select **one** of the two time synchronization daemons; **`chrony (1)`** or **`systemd-timesyncd (2)`** and following the remediation procedure for the selected daemon.

Note: enabling more than one synchronization daemon could lead to unexpected or unreliable results:

1. **`chrony`**

Run the following command to install `chrony`:

```
# apt install chrony
```

Run the following commands to stop and mask the `systemd-timesyncd` daemon:

```
# systemctl stop systemd-timesyncd.service

# systemctl mask systemd-timesyncd.service
```

Note:

- Subsection: ***Configure chrony*** should be followed
- Subsection: ***Configure systemd-timesyncd*** should be skipped

2. **`systemd-timesyncd`**

Run the following command to remove the chrony package:

```
# apt purge chrony
# apt autoremove chrony
```

Note:

- Subsection: ***Configure systemd-timesyncd*** should be followed
- Subsection: ***Configure chrony*** should be skipped

Internal Only - General

References:

1. GRID: NIX-00000154
2. NIST SP 800-53 Rev. 5: AU-3, AU-12

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.4 Standardize Time Synchronization Standardize time synchronization. Configure at least two synchronized time sources across enterprise assets, where supported.		●	●
v7	6.1 Utilize Three Synchronized Time Sources 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.		●	●

Internal Only - General

2.3.2 Configure systemd-timesyncd

`systemd-timesyncd` is a daemon that has been added for synchronizing the system clock across the network. It implements an SNTP client. In contrast to NTP implementations such as chrony or the NTP reference server this only implements a client side, and does not bother with the full NTP complexity, focusing only on querying time from one remote server and synchronizing the local clock to it. The daemon runs with minimal privileges, and has been hooked up with networkd to only operate when network connectivity is available. The daemon saves the current clock to disk every time a new NTP sync has been acquired, and uses this to possibly correct the system clock early at bootup, in order to accommodate for systems that lack an RTC such as the Raspberry Pi and embedded devices, and make sure that time monotonically progresses on these systems, even if it is not always correct. To make use of this daemon a new system user and group "systemd-timesync" needs to be created on installation of systemd.

The default configuration is set during compilation, so configuration is only needed when it is necessary to deviate from those defaults. Initially, the main configuration file in /etc/systemd/ contains commented out entries showing the defaults as a guide to the administrator. Local overrides can be created by editing this file or by creating drop-ins, as described below. Using drop-ins for local configuration is recommended over modifications to the main configuration file.

In addition to the "main" configuration file, drop-in configuration snippets are read from `/usr/lib/systemd/*.conf.d/`, `/usr/local/lib/systemd/*.conf.d/`, and `/etc/systemd/*.conf.d/`. Those drop-ins have higher precedence and override the main configuration file. Files in the *.conf.d/ configuration subdirectories are sorted by their filename in lexicographic order, regardless of in which of the subdirectories they reside. When multiple files specify the same option, for options which accept just a single value, the entry in the file sorted last takes precedence, and for options which accept a list of values, entries are collected as they occur in the sorted files.

When packages need to customize the configuration, they can install drop-ins under /usr/. Files in /etc/ are reserved for the local administrator, who may use this logic to override the configuration files installed by vendor packages. Drop-ins have to be used to override package drop-ins, since the main configuration file has lower precedence. It is recommended to prefix all filenames in those subdirectories with a two-digit number and a dash, to simplify the ordering of the files.

To disable a configuration file supplied by the vendor, the recommended way is to place a symlink to /dev/null in the configuration directory in /etc/, with the same filename as the vendor configuration file.

Note:

- The recommendations in this section only apply if `timesyncd` is in use on the system
- The `systemd-timesyncd` service specifically implements only SNTP.

Internal Only - General

    - This minimalistic service will set the system clock for large offsets or slowly adjust it for smaller deltas
    - More complex use cases are not covered by `systemd-timesyncd`
- **If `chrony` is used, `systemd-timesyncd` should be stopped and masked, and this section skipped**
- **One, and only one, time synchronization method should be in use on the system**

Internal Only - General

2.3.2.1 Ensure systemd-timesyncd configured with authorized timeserver (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`NTP=`

- A space-separated list of NTP server host names or IP addresses. During runtime this list is combined with any per-interface NTP servers acquired from systemd-networkd.service(8). systemd-timesyncd will contact all configured system or per-interface servers in turn, until one responds. When the empty string is assigned, the list of NTP servers is reset, and all prior assignments will have no effect. This setting defaults to an empty list.

`FallbackNTP=`

- A space-separated list of NTP server host names or IP addresses to be used as the fallback NTP servers. Any per-interface NTP servers obtained from systemd-networkd.service(8) take precedence over this setting, as do any servers set via NTP= above. This setting is hence only relevant if no other NTP server information is known. When the empty string is assigned, the list of NTP servers is reset, and all prior assignments will have no effect. If this option is not given, a compiled-in list of NTP servers is used.

Rationale:

Time synchronization is important to support time sensitive security mechanisms and to ensure log files have consistent time records across the enterprise to aid in forensic investigations

Audit:

Run the following command to verify the `NTP` **and/or** `FallbackNTP` option is set to local site approved authoritative time server(s):

Internal Only - General

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

{
   a_output=() a_output2=() a_output3=() a_out=() a_out2=()
a_parlist=("NTP=[^#\n\r]+" "FallbackNTP=[^#\n\r]+")
   l_analyze_cmd="$(readlink -f /bin/systemd-analyze)"
l_systemd_config_file="/etc/systemd/timesyncd.conf"
   f_config_file_parameter_chk()
   {
      l_used_parameter_setting=""
      while IFS= read -r l_file; do
         l_file="$(tr -d '# ' <<< "$l_file")"
         l_used_parameter_setting="$(grep -PHs --
'^\h*'"$l_parameter_name"'\b' "$l_file" | tail -n 1)"
         [ -n "$l_used_parameter_setting" ] && break
      done < <($l_analyze_cmd cat-config "$l_systemd_config_file" | tac |
grep -Pio '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
      if [ -n "$l_used_parameter_setting" ]; then
         while IFS=: read -r l_file_name l_file_parameter; do
            while IFS="=" read -r l_file_parameter_name
l_file_parameter_value; do
               if grep -Pq -- "$l_parameter_value" <<<
"$l_file_parameter_value"; then
                  a_out+=("  - Parameter: \"${l_file_parameter_name// /}\"" \
                  "    correctly set to: \"${l_file_parameter_value// /}\"" \
                  "    in the file: \"$l_file_name\"")
               else
                  a_out2+=("  - Parameter: \"${l_file_parameter_name// /}\""
\
                  "    incorrectly set to: \"${l_file_parameter_value// /}\""
\
                  "    in the file: \"$l_file_name\"" \
                  "    Should be set to: \"$l_value_out\"")
               fi
            done <<< "$l_file_parameter"
         done <<< "$l_used_parameter_setting"
      else
         a_out2+=("  - Parameter: \"$l_parameter_name\" is not set in an
included file" \
         "   *** Note: \"$l_parameter_name\" May be set in a file that's
ignored by load procedure ***")
      fi
   }
   while IFS="=" read -r l_parameter_name l_parameter_value; do # Assess and
check parameters
      l_parameter_name="${l_parameter_name// /}";
l_parameter_value="${l_parameter_value// /}"
      l_value_out="${l_parameter_value//-/ through }";
l_value_out="${l_value_out//|/ or }"
      l_value_out="$(tr -d '(){}' <<< "$l_value_out")"
      f_config_file_parameter_chk
   done < <(printf '%s\n' "${a_parlist[@]}")
   if [ "${#a_out[@]}" -gt 0 ]; then
      a_output+=("${a_out[@]}"); [ "${#a_out2[@]}" -gt 0 ] && a_output3+=("
** INFO: **" "${a_out2[@]}")
   else
      a_output2+=("${a_out2[@]}")
```


Internal Only - General

```
   fi
   if [ "${#a_output2[@]}" -le 0 ]; then
      printf '%s\n' "" "- Audit Result:" "  ** PASS **" "${a_output[@]}" ""
      [ "${#a_output3[@]}" -gt 0 ] && printf '%s\n' "${a_output3[@]}"
   else
      printf '%s\n' "" "- Audit Result:" "  ** FAIL **" " - Reason(s) for
audit failure:" "${a_output2[@]}"
      [ "${#a_output[@]}" -gt 0 ] && printf '%s\n' "" "- Correctly set:"
"${a_output[@]}" ""
   fi
}
```

Note: Please ensure the output for `NTP` and/or `FallbackNTP` is in accordance with local site policy. The timeservers in the example output are provided as an example of possible timeservers and they may not follow local site policy.

Remediation:

Set `NTP` and/or `FallbackNTP` parameters to local site approved authoritative time server(s) in `/etc/systemd/timesyncd.conf` or a file in `/etc/systemd/timesyncd.conf.d/` ending in `.conf` in the `[Time]` section:

Example file:

```
[Time]
NTP=time.nist.gov # Uses the generic name for NIST's time servers
FallbackNTP=time-a-g.nist.gov time-b-g.nist.gov time-c-g.nist.gov # Space
separated list of NIST time servers
```

Example script to create systemd drop-in configuration file:

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

{
   a_settings=("NTP=time.nist.gov" "FallbackNTP=time-a-g.nist.gov time-b-
g.nist.gov time-c-g.nist.gov")
   [ ! -d /etc/systemd/timesyncd.conf.d/ ] && mkdir
/etc/systemd/timesyncd.conf.d/
   if grep -Psq -- '^\h*\[Time\]' /etc/systemd/timesyncd.conf.d/60-
timesyncd.conf; then
      printf '%s\n' "" "${a_settings[@]}" >>
/etc/systemd/timesyncd.conf.d/60-timesyncd.conf
   else
      printf '%s\n' "" "[Time]" "${a_settings[@]}" >>
/etc/systemd/timesyncd.conf.d/60-timesyncd.conf
   fi
}
```

Note: If this setting appears in a canonically later file, or later in the same file, the setting will be overwritten

Run to following command to update the parameters in the service:

Internal Only - General

```
# systemctl reload-or-restart systemd-timesyncd
```

Default Value:

#NTP=

#FallbackNTP=

References:

1. GRID: NIX-00000159
2. NIST SP 800-53 Rev. 5: AU-7, AU-8
3. https://tf.nist.gov/tf-cgi/servers.cgi
4. https://www.freedesktop.org/software/systemd/man/timesyncd.conf.html

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.4 Standardize Time Synchronization Standardize time synchronization. Configure at least two synchronized time sources across enterprise assets, where supported.		●	●
v7	6.1 Utilize Three Synchronized Time Sources 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.		●	●

Internal Only - General

2.3.2.2 Ensure systemd-timesyncd is enabled and running (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

systemd-timesyncd is a daemon that has been added for synchronizing the system clock across the network

Rationale:

systemd-timesyncd needs to be enabled and running in order to synchronize the system to a timeserver.

Time synchronization is important to support time sensitive security mechanisms and to ensure log files have consistent time records across the enterprise to aid in forensic investigations

Audit:

- IF - systemd-timesyncd is in use on the system, run the following commands:

Run the following command to verify that the `systemd-timesyncd` service is enabled:

```
# systemctl is-enabled systemd-timesyncd.service

enabled
```

Run the following command to verify that the `systemd-timesyncd` service is active:

```
# systemctl is-active systemd-timesyncd.service

active
```

Remediation:

- IF - `systemd-timesyncd` is in use on the system, run the following commands:

Run the following command to unmask `systemd-timesyncd.service`:

```
# systemctl unmask systemd-timesyncd.service
```

Run the following command to enable and start `systemd-timesyncd.service`:

```
# systemctl --now enable systemd-timesyncd.service
```

- OR -

Internal Only - General

If another time synchronization service is in use on the system, run the following command to stop and mask `systemd-timesyncd`:

```
# systemctl --now mask systemd-timesyncd.service
```

References:

1. GRID: NIX-00000160
2. NIST SP 800-53 Rev. 5: AU-7, AU-8

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.4 Standardize Time Synchronization Standardize time synchronization. Configure at least two synchronized time sources across enterprise assets, where supported.		●	●
v7	6.1 Utilize Three Synchronized Time Sources 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.		●	●

Internal Only - General

2.3.3 Configure chrony

`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.

`chrony` can be configured to be a client and/or a server.

More information on `chrony` can be found at: http://chrony.tuxfamily.org/.

Note:

- If `systemd-timesyncd` is being used, `chrony` should be removed and this section skipped
- Only one time synchronization method should be in use on the system

Internal Only - General

2.3.3.1 Ensure chrony is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

- server
  - The server directive specifies an NTP server which can be used as a time source. The client-server relationship is strictly hierarchical: a client might synchronize its system time to that of the server, but the server's system time will never be influenced by that of a client.
  - This directive can be used multiple times to specify multiple servers.
  - The directive is immediately followed by either the name of the server, or its IP address.
- pool
  - The syntax of this directive is similar to that for the server directive, except that it is used to specify a pool of NTP servers rather than a single NTP server. The pool name is expected to resolve to multiple addresses which might change over time.
  - This directive can be used multiple times to specify multiple pools.
  - All options valid in the server directive can be used in this directive too.

Rationale:

Time synchronization is important to support time sensitive security mechanisms and to ensure log files have consistent time records across the enterprise to aid in forensic investigations

Audit:

- IF - `chrony` is in use on the system, run the following script to ensure `chrony` is configured with an authorized timeserver:

Internal Only - General

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

{
   a_output=() a_output2=() a_config_files=("/etc/chrony/chrony.conf")
   l_include='(confdir|sourcedir)' l_parameter_name='(server|pool)'
l_parameter_value='.+'
   while IFS= read -r l_conf_loc; do
      l_dir="" l_ext=""
      if [ -d "$l_conf_loc" ]; then
         l_dir="$l_conf_loc" l_ext="*"
      elif  grep -Psq '\/\*\.([^#/\n\r]+)?\h*$' <<< "$l_conf_loc" || [ -f
"$(readlink -f "$l_conf_loc")" ]; then
         l_dir="$(dirname "$l_conf_loc")" l_ext="$(basename "$l_conf_loc")"
      fi
      if [[ -n "$l_dir" && -n "$l_ext" ]]; then
         while IFS= read -r -d $'\0' l_file_name; do
            [ -f "$(readlink -f "$l_file_name")" ] &&
a_config_files+=("$(readlink -f "$l_file_name")")
         done < <(find -L "$l_dir" -type f -name "$l_ext" -print0
2>/dev/null)
      fi
   done < <(awk '$1~/^\s*'"$l_include"'$/{print $2}' "${a_config_files[*]}"
2>/dev/null)
   for l_file in "${a_config_files[@]}"; do
      l_parameter_line="$(grep -Psi
'^\h*'"$l_parameter_name"'(\h+|\h*:\h*)'"$l_parameter_value"'\b' "$l_file")"
      [ -n "$l_parameter_line" ] && a_output+=("  - Parameter: \"$(tr -d '()'
<<< ${l_parameter_name//|/ or })\"" \
      "    Exists in the file: \"$l_file\" as:" "$l_parameter_line")
   done
   [ "${#a_output[@]}" -le "0" ] && a_output2+=("  - Parameter: \"$(tr -d
'()' <<< ${l_parameter_name//|/ or })\"" \
   "    Does not exist in the chrony configuration")
   if [ "${#a_output2[@]}" -le 0 ]; then
      printf '%s\n' "" "- Audit Result:" "  ** PASS **" "${a_output[@]}" ""
   else
      printf '%s\n' "" "- Audit Result:" "  ** FAIL **" " - Reason(s) for
audit failure:" "${a_output2[@]}"
   fi
}
```

Remediation:

Edit `/etc/chrony/chrony.conf` or a file ending in `.sources` in `/etc/chrony/sources.d/` and add or edit server or pool lines as appropriate according to local site policy:

Edit the `Chrony` configuration and add or edit the server and/or pool lines returned by the Audit Procedure as appropriate according to local site policy

```
<[server|pool]> <[remote-server|remote-pool]>
```

Example script to add a drop-in configuration for the `pool` directive:

Internal Only - General

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

{
   [ ! -d "/etc/chrony/sources.d/" ] && mkdir /etc/chrony/sources.d/
   printf '%s\n' "" "#The maxsources option is unique to the pool directive"
\
   "pool time.nist.gov iburst maxsources 4" >> /etc/chrony/sources.d/60-
sources.sources
   chronyc reload sources &>/dev/null
}
```

Example script to add a drop-in configuration for the `server` directive:

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

{
   [ ! -d "/etc/chrony/sources.d/" ] && mkdir /etc/chrony/sources.d/
   printf '%s\n' "" "server time-a-g.nist.gov iburst" "server 132.163.97.3
iburst" \
   "server time-d-b.nist.gov iburst" >> /etc/chrony/sources.d/60-
sources.sources
   chronyc reload sources &>/dev/null
}
```

Run the following command to reload the `chronyd` config:

```
# systemctl reload-or-restart chronyd
```

References:

1. GRID: NIX-00000155
2. NIST SP 800-53 Rev. 5: AU-3, AU-12
3. chrony.conf(5) Manual Page
4. https://tf.nist.gov/tf-cgi/servers.cgi

Additional Information:

If pool and/or server directive(s) are set in a sources file in `/etc/chrony/sources.d`, the line:

```
sourcedir /etc/chrony/sources.d
```

must be present in `/etc/chrony/chrony.conf`

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.4 Standardize Time Synchronization Standardize time synchronization. Configure at least two synchronized time sources across enterprise assets, where supported.		●	●

Internal Only - General

Controls Version	Control	IG 1	IG 2	IG 3
v7	6.1 Utilize Three Synchronized Time Sources 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.		●	●

Internal Only - General

2.3.3.2 Ensure chrony is running as user _chrony (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `chrony` package is installed with a dedicated user account `_chrony`. This account is granted the access required by the `chronyd` service

Rationale:

The `chronyd` service should run with only the required privlidges

Audit:

- IF - `chrony` is in use on the system, run the following command to verify the `chronyd` service is being run as the `_chrony` user:

```
# ps -ef | awk '(/[c]hronyd/ && $1!="_chrony") { print $1 }'
```

Nothing should be returned

Remediation:

Add or edit the `user` line to `/etc/chrony/chrony.conf` or a file ending in `.conf` in `/etc/chrony/conf.d/`:

```
user _chrony
```

- OR -

If another time synchronization service is in use on the system, run the following command to remove `chrony` from the system:

```
# apt purge chrony
# apt autoremove chrony
```

Default Value:

user _chrony

References:

1. GRID: NIX-00000158
2. NIST SP 800-53 Rev. 5: AU-8

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

2.3.3.3 Ensure chrony is enabled and running (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

chrony is a daemon for synchronizing the system clock across the network

Rationale:

chrony needs to be enabled and running in order to synchronize the system to a timeserver.

Time synchronization is important to support time sensitive security mechanisms and to ensure log files have consistent time records across the enterprise to aid in forensic investigations

Audit:

- IF - chrony is in use on the system, run the following commands:

Run the following command to verify that the `chrony` service is enabled:

```
# systemctl is-enabled chrony.service

enabled
```

Run the following command to verify that the `chrony` service is active:

```
# systemctl is-active chrony.service

active
```

Remediation:

- IF - `chrony` is in use on the system, run the following commands:

Run the following command to unmask `chrony.service`:

```
# systemctl unmask chrony.service
```

Run the following command to enable and start `chrony.service`:

```
# systemctl --now enable chrony.service
```

- OR -

If another time synchronization service is in use on the system, run the following command to remove `chrony`:

Internal Only - General

```
# apt purge chrony
# apt autoremove chrony
```

References:

1. GRID: NIX-00000156
2. NIST SP 800-53 Rev. 5: AU-8

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.4 Standardize Time Synchronization Standardize time synchronization. Configure at least two synchronized time sources across enterprise assets, where supported.		●	●
v7	6.1 Utilize Three Synchronized Time Sources 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.		●	●

2.4 Job Schedulers

A job scheduler is used to execute jobs, commands, or shell scripts, at fixed times, dates, or intervals

Internal Only - General

2.4.1 Configure cron

`cron` is a time based job scheduler

- IF - `cron` is not installed on the system, this sub section can be skipped

Note: Other methods such as `systemd timers` exist for scheduling jobs. If an alternate method is in use, it should be secured in accordance with local site policy

Internal Only - General

2.4.1.1 Ensure cron daemon is enabled and active (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:

- IF - cron is installed on the system:

Run the following command to verify `cron` is enabled:

```
# systemctl list-unit-files | awk '$1~/^crond?\.service/{print $2}'

enabled
```

Run the following command to verify that `cron` is active:

```
# systemctl list-units | awk '$1~/^crond?\.service/{print $3}'

active
```

Remediation:

- IF - cron is installed on the system:

Run the following commands to unmask, enable, and start `cron`:

```
# systemctl unmask "$(systemctl list-unit-files | awk
'$1~/^crond?\.service/{print $1}')"
# systemctl --now enable "$(systemctl list-unit-files | awk
'$1~/^crond?\.service/{print $1}')"
```

References:

1. GRID: NIX-00000161
2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

2.4.1.2 Ensure access to /etc/crontab is 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:

- IF - cron is installed on the system:

Run the following command and verify `Uid` and `Gid` are both `0/root` and `Access` does not grant permissions to `group` or `other` :

```
# stat -Lc 'Access: (%a/%A) Uid: ( %u/ %U) Gid: ( %g/ %G)' /etc/crontab

Access: (600/-rw-------) Uid: ( 0/ root) Gid: ( 0/ root)
```

Remediation:

- IF - cron is installed on the system:

Run the following commands to set ownership and permissions on `/etc/crontab`:

```
# chown root:root /etc/crontab
# chmod og-rwx /etc/crontab
```

Default Value:

Access: (644/-rw-r--r--) Uid: (0/ root) Gid: (0/ root)

References:

1. GRID: NIX-00000162
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

Internal Only - General

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.	●	●	●

Internal Only - General

2.4.1.3 Ensure access to /etc/cron.hourly is 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:

- IF - cron is installed on the system:

Run the following command and verify `Uid` and `Gid` are both `0/root` and `Access` does not grant permissions to `group` or `other`:

```
# stat -Lc 'Access: (%a/%A) Uid: ( %u/ %U) Gid: ( %g/ %G)' /etc/cron.hourly/

Access: (700/drwx------) Uid: ( 0/ root) Gid: ( 0/ root)
```

Remediation:

- IF - cron is installed on the system:

Run the following commands to set ownership and permissions on the `/etc/cron.hourly` directory:

```
# chown root:root /etc/cron.hourly/
# chmod og-rwx /etc/cron.hourly/
```

Default Value:

Access: (755/drwxr-xr-x) Uid: (0/ root) Gid: (0/ root)

Internal Only - General

References:

1. GRID: NIX-00000163
2. NIST SP 800-53 Rev. 5: AC-3. MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 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.	●	●	●

Internal Only - General

2.4.1.4 Ensure access to /etc/cron.daily is 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:

- IF - cron is installed on the system:

Run the following command and verify `Uid` and `Gid` are both `0/root` and `Access` does not grant permissions to `group` or `other`:

```
# stat -Lc 'Access: (%a/%A) Uid: ( %u/ %U) Gid: ( %g/ %G)' /etc/cron.daily/

Access: (700/drwx------) Uid: ( 0/ root) Gid: ( 0/ root)
```

Remediation:

- IF - cron is installed on the system:

Run the following commands to set ownership and permissions on the `/etc/cron.daily` directory:

```
# chown root:root /etc/cron.daily/
# chmod og-rwx /etc/cron.daily/
```

Default Value:

Access: (755/drwxr-xr-x) Uid: (0/ root) Gid: (0/ root)

References:

1. GRID: NIX-00000164
2. NIST SP 800-53 Rev. 5: AC-3. MP-2

Internal Only - General

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.	●	●	●

Internal Only - General

2.4.1.5 Ensure access to /etc/cron.weekly is 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:

- IF - cron is installed on the system:

Run the following command and verify `Uid` and `Gid` are both `0/root` and `Access` does not grant permissions to `group` or `other`:

```
# stat -Lc 'Access: (%a/%A) Uid: ( %u/ %U) Gid: ( %g/ %G)' /etc/cron.weekly/

Access: (700/drwx------) Uid: ( 0/ root) Gid: ( 0/ root)
```

Remediation:

- IF - cron is installed on the system:

Run the following commands to set ownership and permissions on the `/etc/cron.weekly` directory:

```
# chown root:root /etc/cron.weekly/
# chmod og-rwx /etc/cron.weekly/
```

Default Value:

Access: (755/drwxr-xr-x) Uid: (0/ root) Gid: (0/ root)

Internal Only - General

References:

1. GRID: NIX-00000165
2. NIST SP 800-53 Rev. 5: AC-3. MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 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.	●	●	●

Internal Only - General

2.4.1.6 Ensure access to /etc/cron.monthly is 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:

- IF - cron is installed on the system:

Run the following command and verify `Uid` and `Gid` are both `0/root` and `Access` does not grant permissions to `group` or `other`:

```
# stat -Lc 'Access: (%a/%A) Uid: ( %u/ %U) Gid: ( %g/ %G)' /etc/cron.monthly/

Access: (700/drwx------) Uid: ( 0/ root) Gid: ( 0/ root)
```

Remediation:

- IF - cron is installed on the system:

Run the following commands to set ownership and permissions on the `/etc/cron.monthly` directory:

```
# chown root:root /etc/cron.monthly/
# chmod og-rwx /etc/cron.monthly/
```

Default Value:

Access: (755/drwxr-xr-x) Uid: (0/ root) Gid: (0/ root)

Internal Only - General

References:

1. GRID: NIX-00000166
2. NIST SP 800-53 Rev. 5: AC-3. MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 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.	●	●	●

Internal Only - General

2.4.1.7 Ensure access to /etc/cron.yearly is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `/etc/cron.yearly` directory contains system cron jobs that run on an annual basis. Files in this directory are edited directly by system administrators rather than through the `crontab` command. The commands below restrict read, write, and search access to `root` only, preventing regular users from accessing this directory.

Rationale:

Granting write access to `/etc/cron.yearly` for non-privileged users would allow them to introduce arbitrary commands that run annually as root, providing a path to unauthorized privilege escalation. Granting read access would expose the contents of scheduled jobs, potentially revealing sensitive system information or attack opportunities.

Audit:

- IF - cron is installed on the system:

Run the following command and verify `Uid` and `Gid` are both `0/root` and `Access` does not grant permissions to `group` or `other`:

```
# stat -Lc 'Access: (%a/%A) Uid: ( %u/ %U) Gid: ( %g/ %G)' /etc/cron.yearly/

Access: (700/drwx------) Uid: ( 0/ root) Gid: ( 0/ root)
```

Note: If the directory does not exist, this is considered a passing state.

Remediation:

- IF - cron is installed on the system:

Run the following commands to set ownership and permissions on the `/etc/cron.yearly` directory:

```
# chown root:root /etc/cron.yearly/
# chmod og-rwx /etc/cron.yearly/
```

Default Value:

Access: (755/drwxr-xr-x) Uid: (0/ root) Gid: (0/ root)

Internal Only - General

References:

1. GRID: NIX-00000167
2. NIST SP 800-53 Rev. 5: AC-3. MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 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.	●	●	●

Internal Only - General

2.4.1.8 Ensure access to /etc/cron.d is 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:

- IF - cron is installed on the system:

Run the following command and verify `Uid` and `Gid` are both `0/root` and `Access` does not grant permissions to `group` or `other`:

```
# stat -Lc 'Access: (%a/%A) Uid: ( %u/ %U) Gid: ( %g/ %G)' /etc/cron.d/

Access: (700/drwx------) Uid: ( 0/ root) Gid: ( 0/ root)
```

Remediation:

- IF - cron is installed on the system:

Run the following commands to set ownership and permissions on the `/etc/cron.d` directory:

```
# chown root:root /etc/cron.d/
# chmod og-rwx /etc/cron.d/
```

Default Value:

Access: (755/drwxr-xr-x) Uid: (0/ root) Gid: (0/ root)

References:

1. GRID: NIX-00000168

Internal Only - General

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

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 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.	●	●	●

Internal Only - General

2.4.1.9 Ensure access to crontab is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`crontab` is the program used to install, deinstall, or list the tables used to drive the cron daemon. Each user can have their own crontab, and though files are created, they are not intended to be edited directly.

If the `/etc/cron.allow` file exists, then you must be listed (one user per line) therein in order to be allowed to use this command. If the `/etc/cron.allow` file does not exist but the `/etc/cron.deny` file does exist, then you must not be listed in the `/etc/cron.deny` file in order to use this command.

If neither of these files exists, then depending on site-dependent configuration parameters, only the super user will be allowed to use this command, or all users will be able to use this command.

If both files exist then `/etc/cron.allow` takes precedence. Which means that `/etc/cron.deny` is not considered and your user must be listed in `/etc/cron.allow` in order to be able to use the crontab.

Regardless of the existence of any of these files, the root administrative user is always allowed to setup a crontab.

The files `/etc/cron.allow` and `/etc/cron.deny`, if they exist, must be either world-readable, or readable by group `crontab`. If they are not, then cron will deny access to all users until the permissions are fixed.

There is one file for each user's crontab. Users are not allowed to edit the file directly to ensure that only users allowed by the system to run periodic tasks can add them, and only syntactically correct crontabs will be written. This is enforced by having the directory writable only by the `crontab` group and configuring crontab command with the setgid bid set for that specific group.

Note:

- The files `/etc/cron.allow` and `/etc/cron.deny`, if they exist, only controls administrative access to the crontab command for scheduling and modifying cron jobs
- If you are using SELinux or AppArmor, these security modules may prevent crontab access. Check the audit logs for any related `permission denied` entries and adjust your policies if necessary.

Internal Only - General

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.

Impact:

When the `/etc/cron.allow` file exists, a user must be listed, one user per line, in the `/etc/cron.allow` file in order to be allowed to use the `crontab` command.

Audit:

- IF - cron is installed on the system:

Run the following command to verify `/etc/cron.allow`:

- Exists
- Is mode `0640` or more restrictive
- Is owned by the user `root`
- Is group owned by the group `root` **- OR -** the group `crontab`

```
# stat -Lc 'Access: (%a/%A) Owner: (%U) Group: (%G)' /etc/cron.allow
```

Verify the returned value is:

```
Access: (640/-rw-r-----) Owner: (root) Group: (root)
- OR -
Access: (640/-rw-r-----) Owner: (root) Group: (crontab)
```

Run the following command to verify either `cron.deny` doesn't exist or is:

- Mode `0640` or more restrictive
- Owned by the user `root`
- Is group owned by the group `root` **- OR -** the group `crontab`

```
# [ -e "/etc/cron.deny" ] && stat -Lc 'Access: (%a/%A) Owner: (%U) Group:
(%G)' /etc/cron.deny
```

Verify either nothing is returned **- OR -** returned value is one of the following:

```
Access: (640/-rw-r-----) Owner: (root) Group: (root)
- OR -
Access: (640/-rw-r-----) Owner: (root) Group: (crontab)
```

Note: On systems where cron is configured to use the group `crontab`, if the group `crontab` is not set as the owner of `cron.allow`, then cron will deny access to all users and you will see an error similar to:

Internal Only - General

```
You (<USERNAME>) are not allowed to use this program (crontab)
See crontab(1) for more information
```

Remediation:

- IF - cron is installed on the system:

Run the following script to:

- Create `/etc/cron.allow` if it doesn't exist
- Change owner to user `root`
- Change group owner to group `root` **- OR -** group `crontab` if it exists
- Change mode to `640` or more restrictive

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

{
   [ ! -e "/etc/cron.allow" ] && touch /etc/cron.allow
   chmod u-x,g-wx,o-rwx /etc/cron.allow
   if grep -Pq -- '^\h*crontab\:' /etc/group; then
      chown root:crontab /etc/cron.allow
   else
      chown root:root /etc/cron.allow
   fi
}
```

- IF - `/etc/cron.deny` exists, run the following script to:

- Change owner to user `root`
- Change group owner to group `root` **- OR -** group `crontab` if it exists
- Change mode to `640` or more restrictive

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

{
   if [ -e "/etc/cron.deny" ]; then
      chmod u-x,g-wx,o-rwx /etc/cron.deny
      if grep -Pq -- '^\h*crontab\:' /etc/group; then
         chown root:crontab /etc/cron.deny
      else
         chown root:root /etc/cron.deny
      fi
   fi
}
```

Note:

- When the `/etc/cron.allow` file exists, **a user must be listed, one user per line, in the `/etc/cron.allow` file in order to be allowed to use the `crontab` command**.

Internal Only - General

- If both `/etc/cron.allow` and `/etc/cron.deny` exist, `/etc/cron.allow` takes precedence. Which means that `/etc/cron.deny` is not considered and users must be listed in `/etc/cron.allow` in order to use the crontab.
- If the `/etc/cron.allow` file does not exist but the `/etc/cron.deny` file does exist, then you must not be listed in the `/etc/cron.deny` file in order to use the `crontab` command.
- Regardless of the existence of any of these files, the root administrative user is always allowed to setup a crontab.
- On systems where cron is configured to use the group `crontab`, if the group `crontab` is not set as the owner of `/etc/cron.allow`, then cron will deny access to all users but `root`, and you will see an error similar to:

```
You (<USERNAME>) are not allowed to use this program (crontab)
See crontab(1) for more information
```

References:

1. GRID: NIX-00000169
2. NIST SP 800-53 Rev. 5: AC-3. MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 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.	●	●	●

Internal Only - General

2.4.2 Configure at

`at` is a command-line utility used to schedule a job for later execution

Note: if `at` is not installed on the system, this section can be skipped

Internal Only - General

2.4.2.1 Ensure access to at is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`at` allows fairly complex time specifications, extending the POSIX.2 standard. It accepts times of the form HH:MM to run a job at a specific time of day. (If that time is already past, the next day is assumed.) You may also specify midnight, noon, or teatime (4pm) and you can have a time-of-day suffixed with AM or PM for running in the morning or the evening. You can also say what day the job will be run, by giving a date in the form month-name day with an optional year, or giving a date of the form MMDD[CC]YY, MM/DD/[CC]YY, DD.MM.[CC]YY or [CC]YY-MM-DD. The specification of a date must follow the specification of the time of day. You can also give times like now + count time-units, where the time-units can be minutes, hours, days, or weeks and you can tell at to run the job today by suffixing the time with today and to run the job tomorrow by suffixing the time with tomorrow.

The `/etc/at.allow` and `/etc/at.deny` files determine which user can submit commands for later execution via at or batch. The format of the files is a list of usernames, one on each line. Whitespace is not permitted. If the file `/etc/at.allow` exists, only usernames mentioned in it are allowed to use at. If `/etc/at.allow` does not exist, `/etc/at.deny` is checked, every username not mentioned in it is then allowed to use at. An empty `/etc/at.deny` means that every user may use at. If neither file exists, only the superuser is allowed to use at.

Rationale:

On many systems, only the system administrator is authorized to schedule `at` jobs. Using the `at.allow` file to control who can run `at` jobs enforces this policy. It is easier to manage an allow list than a deny list. In a deny list, you could potentially add a user ID to the system and forget to add it to the deny files.

Audit:

- IF - at is installed on the system:

Run the following command to verify `/etc/at.allow`:

- Exists
- Is mode `0640` or more restrictive
- Is owned by the user `root`
- Is group owned by the group `daemon` or group `root`

Internal Only - General

```
# stat -Lc 'Access: (%a/%A) Owner: (%U) Group: (%G)' /etc/at.allow

Access: (640/-rw-r-----) Owner: (root) Group: (daemon)
-OR-
Access: (640/-rw-r-----) Owner: (root) Group: (root)
```

Verify mode is `640` or more restrictive, owner is `root`, and group is `daemon` or `root`

Run the following command to verify `at.deny` doesn't exist, **-OR-** is:

- Mode `0640` or more restrictive
- Owned by the user `root`
- Group owned by the group `daemon` or group `root`

```
# [ -e "/etc/at.deny" ] && stat -Lc 'Access: (%a/%A) Owner: (%U) Group: (%G)'
/etc/at.deny

Access: (640/-rw-r-----) Owner: (root) Group: (daemon)
-OR-
Access: (640/-rw-r-----) Owner: (root) Group: (root)
-OR-
Nothing is returned
```

If a value is returned, verify mode is 640 or more restrictive, owner is `root`, and group is `daemon` or `root`

Remediation:

- IF - at is installed on the system:

Run the following script to:

- `/etc/at.allow`:
  - Create the file if it doesn't exist
  - Change owner or user `root`
  - If group `daemon` exists, change to group `daemon`, else change group to `root`
  - Change mode to `640` or more restrictive
- **- IF -** `/etc/at.deny` exists:
  - Change owner or user `root`
  - If group `daemon` exists, change to group `daemon`, else change group to `root`
  - Change mode to `640` or more restrictive

Internal Only - General

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

{
   grep -Pq -- '^daemon\b' /etc/group && l_group="daemon" || l_group="root"
   [ ! -e "/etc/at.allow" ] && touch /etc/at.allow
   chown root:"$l_group" /etc/at.allow
   chmod u-x,g-wx,o-rwx /etc/at.allow
   [ -e "/etc/at.deny" ] && chown root:"$l_group" /etc/at.deny
   [ -e "/etc/at.deny" ] && chmod u-x,g-wx,o-rwx /etc/at.deny
}
```

References:

1. GRID: NIX-00000170
2. NIST SP 800-53 Rev. 5: AC-3. MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 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.	●	●	●

3 Network

This section provides guidance on for securing the network configuration of the system

Internal Only - General

3.1 Configure Network Devices

To reduce the attack surface of a system, unused devices should be disabled.

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

Internal Only - General

3.1.1 Ensure IPv6 status is identified (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Internet Protocol Version 6 (IPv6) is the most recent version of Internet Protocol (IP). It's designed to supply IP addressing and additional security to support the predicted growth of connected devices. IPv6 is based on 128-bit addressing and can support 340 undecillion, which is 340,282,366,920,938,463,463,374,607,431,768,211,456 unique addresses.

Features of IPv6

- Hierarchical addressing and routing infrastructure
- Statefull and Stateless configuration
- Support for quality of service (QoS)
- An ideal protocol for neighboring node interaction

Rationale:

IETF RFC 4038 recommends that applications are built with an assumption of dual stack. It is recommended that IPv6 be enabled and configured in accordance with Benchmark recommendations.

- IF - dual stack and IPv6 are not used in your environment, IPv6 may be disabled to reduce the attack surface of the system, and recommendations pertaining to IPv6 can be skipped.

Note: It is recommended that IPv6 be enabled and configured unless this is against local site policy

Impact:

IETF RFC 4038 recommends that applications are built with an assumption of dual stack.

When enabled, IPv6 will require additional configuration to reduce risk to the system.

Audit:

Run the following script to identify if IPv6 is enabled on the system:

Internal Only - General

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

{
   l_output=""
   ! grep -Pqs -- '^\h*0\b' /sys/module/ipv6/parameters/disable &&
l_output="- IPv6 is not enabled"
   if sysctl net.ipv6.conf.all.disable_ipv6 | grep -Pqs --
"^\h*net\.ipv6\.conf\.all\.disable_ipv6\h*=\h*1\b" && \
      sysctl net.ipv6.conf.default.disable_ipv6 | grep -Pqs --
"^\h*net\.ipv6\.conf\.default\.disable_ipv6\h*=\h*1\b"; then
      l_output="- IPv6 is not enabled"
   fi
   [ -z "$l_output" ] && l_output="- IPv6 is enabled"
   echo -e "\n$l_output\n"
}
```

Remediation:

Enable or disable IPv6 in accordance with system requirements and local site policy

Default Value:

IPv6 is enabled

References:

1. GRID: NIX-00000172
2. NIST SP 800-53 Rev. 5: CM-7

Additional Information:

Having more addresses has grown in importance with the expansion of smart devices and connectivity. IPv6 provides more than enough globally unique IP addresses for every networked device currently on the planet, helping ensure providers can keep pace with the expected proliferation of IP-based devices.

CIS Controls:

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

Internal Only - General

Internal Only - General

3.1.2 Ensure wireless interfaces are not available (Automated)

Profile Applicability:

- Level 1 - Server

Description:

Wireless networking is used when wired networks are unavailable. Wireless interfaces use kernel modules loaded at runtime.

Rationale:

- IF - wireless is not required, wireless kernel modules should be disabled to reduce the potential attack surface. Wireless interfaces can be used as a secondary network path to bypass network monitoring, exfiltrate data, or introduce rogue access points that circumvent perimeter controls. Servers and workstations that operate in wired-only environments have no operational need for wireless capability.

Impact:

Many laptop workstations and some desktop workstations require wireless connectivity for normal operation. Applying this recommendation to such systems will disable wireless network access, which may disrupt users. Verify that no wireless connectivity is operationally required before applying remediation

Audit:

Run the following script to verify no wireless interfaces are active on the system:

Internal Only - General

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

{
   l_output="" l_output2=""
   module_chk()
   {
      # Check how module will be loaded
      l_loadable="$(modprobe -n -v "$l_mname")"
      if grep -Pq -- '^\h*install \/bin\/(true|false)' <<< "$l_loadable";
then
         l_output="$l_output\n - module: \"$l_mname\" is not loadable:
\"$l_loadable\""
      else
         l_output2="$l_output2\n - module: \"$l_mname\" is loadable:
\"$l_loadable\""
      fi
      # Check is the module currently loaded
      if ! lsmod | grep "$l_mname" > /dev/null 2>&1; then
         l_output="$l_output\n - module: \"$l_mname\" is not loaded"
      else
         l_output2="$l_output2\n - module: \"$l_mname\" is loaded"
      fi
      # Check if the module is deny listed
      if modprobe --showconfig | grep -Pq -- "^\h*blacklist\h+$l_mname\b";
then
         l_output="$l_output\n - module: \"$l_mname\" is deny listed in:
\"$(grep -Pl -- "^\h*blacklist\h+$l_mname\b" /etc/modprobe.d/*)\""
      else
         l_output2="$l_output2\n - module: \"$l_mname\" is not deny listed"
      fi
   }
   if [ -n "$(find /sys/class/net/*/ -type d -name wireless)" ]; then
      l_dname=$(for driverdir in $(find /sys/class/net/*/ -type d -name
wireless | xargs -0 dirname); do basename "$(readlink -f
"$driverdir"/device/driver/module)";done | sort -u)
      for l_mname in $l_dname; do
         module_chk
      done
   fi
   # Report results. If no failures output in l_output2, we pass
   if [ -z "$l_output2" ]; then
      echo -e "\n- Audit Result:\n  ** PASS **"
      if [ -z "$l_output" ]; then
         echo -e "\n - System has no wireless NICs installed"
      else
         echo -e "\n$l_output\n"
      fi
   else
      echo -e "\n- Audit Result:\n  ** FAIL **\n - Reason(s) for audit
failure:\n$l_output2\n"
      [ -n "$l_output" ] && echo -e "\n- Correctly set:\n$l_output\n"
   fi
}
```

Remediation:

Internal Only - General

Run the following command to disable any wireless interfaces:

```
# find /lib/modules/`uname -r`/kernel/drivers/net/wireless -name '*.ko' -
printf 'install %f /bin/false\nblacklist %f\n\n' | sed 's/\.ko//1' >>
/etc/modprobe.d/blacklist-wireless.conf
```

Note: the `*.conf` file in `/etc/modprobe.d/` in the above command can renamed as needed.

After running this command, reload module configurations and verify no wireless modules remain loaded:

```
# modprobe --showconfig | grep -E '(install|blacklist).*80211|wifi|wireless'
# lsmod | grep -E '80211|wifi|wireless'
```

Nothing should be returned by the `lsmod` command.

References:

1. GRID: NIX-00000173
2. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●
v7	15.4 Disable Wireless Access on Devices if Not Required Disable wireless access on devices that do not have a business purpose for wireless access.			●
v7	15.5 Limit Wireless Access on Client Devices 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.			●

Internal Only - General

3.1.3 Ensure bluetooth services are not in use (Automated)

Profile Applicability:

- Level 1 - Server
- Level 2 - Workstation

Description:

Bluetooth is a short-range wireless technology standard that is used for exchanging data between devices over short distances. It employs UHF radio waves in the ISM bands, from 2.402 GHz to 2.48 GHz. It is mainly used as an alternative to wire connections.

Rationale:

An attacker may be able to find a way to access or corrupt your data. One example of this type of activity is `bluesnarfing`, which refers to attackers using a Bluetooth connection to steal information off of your Bluetooth device. Also, viruses or other malicious code can take advantage of Bluetooth technology to infect other devices. If you are infected, your data may be corrupted, compromised, stolen, or lost.

Impact:

Many personal electronic devices (PEDs) use Bluetooth technology. For example, you may be able to operate your computer with a wireless keyboard. Disabling Bluetooth will prevent these devices from connecting to the system.

There may be packages that are dependent on the `bluez` package. If the `bluez` package is removed, these dependent packages will be removed as well. Before removing the `bluez` package, review any dependent packages to determine if they are required on the system.

-IF- a dependent package is required: stop and mask `bluetooth.service` leaving the `bluez` package installed.

Audit:

Run the following command to verify the `bluez` package is not installed:

```
# dpkg-query -s bluez &>/dev/null && echo "bluez is installed"
```

Nothing should be returned.

- OR -

- IF - the `bluez` package is required as a dependency:

Run the following command to verify `bluetooth.service` is not enabled:

Internal Only - General

```
# systemctl is-enabled bluetooth.service 2>/dev/null | grep 'enabled'
```

Nothing should be returned.

Run the following command to verify `bluetooth.service` is not active:

```
# systemctl is-active bluetooth.service 2>/dev/null | grep '^active'
```

Nothing should be returned.

Note: If the package is required for a dependency

- Ensure the dependent package is approved by local site policy
- Ensure stopping and masking the service and/or socket meets local site policy

Remediation:

Run the following commands to stop `bluetooth.service`, and remove the `bluez` package:

```
# systemctl stop bluetooth.service
# apt purge bluez
```

- OR -

- IF - the `bluez` package is required as a dependency:

Run the following commands to stop and mask `bluetooth.service`:

```
# systemctl stop bluetooth.service
# systemctl mask bluetooth.service
```

Note: A reboot may be required

References:

1. GRID: NIX-00000171
2. https://www.cisa.gov/tips/st05-015
3. NIST SP 800-53 Rev. 5: CM-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●

Internal Only - General

Controls Version	Control	IG 1	IG 2	IG 3
v7	9.2 Ensure Only Approved Ports, Protocols and Services Are Running Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.		●	●

Internal Only - General

3.2 Configure Network Kernel Modules

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.

Internal Only - General

3.2.1 Ensure atm kernel module is not available (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The Asynchronous Transfer Mode (ATM) is a protocol operating on network, data link, and physical layers, based on virtual circuits and virtual paths. Disabling ATM protects the system against exploitation of any laws in its implementation.

Rationale:

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

Audit:

Verify the `atm` kernel module is not available on the system or has been disabled.

1. Run the following script to determine if the `atm` kernel module is available on the system:

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

{
   l_mod_name="atm" l_mod_type="net"
   while IFS= read -r l_mod_path; do
      if [ -d "$l_mod_path/${l_mod_name//-/\/}" ] &&  \
      [ -n "$(ls -A "$l_mod_path/${l_mod_name//-/\/}")" ]; then
         printf '%s\n' "$l_mod_name exists in $l_mod_path"
      fi
   done < <(readlink -e /usr/lib/modules/**/kernel/$l_mod_type \
   || readlink -e /lib/modules/**/kernel/$l_mod_type)
}
```

If nothing is returned, the `atm` kernel module is not available on the system and no further audit steps are required.

Note: Some systems may include the `atm` filesystem as part of the kernel opposed to being available as a kernel module. In this case, the above audit will not return anything. This is also considered a passing state.

If anything is returned by the above script:

2. verify the `atm` kernel module is not loaded and not loadable by performing the following:

Run the following command to verify the `atm` kernel module is not loaded:

Internal Only - General

```
# lsmod | grep 'atm'
```

Nothing should be returned

Run the following command to verify the `atm` kernel module is not loadable:

```
# modprobe --showconfig | grep -P -- '\b(install|blacklist)\h+atm\b'
```

Verify the output includes:

```
blacklist atm
 -AND-
install atm /bin/false
 -OR-
install atm /bin/true
```

Example output:

```
blacklist atm
install atm /bin/false
```

Remediation:

Unload and disable the `atm` kernel module.

1. Run the following commands to unload the `atm` kernel module:

```
# modprobe -r atm 2>/dev/null
# rmmod can 2>/dev/null
```

2. Perform the following to disable the `atm` kernel module:

Create a file ending in `.conf` with `install can /bin/false` in the `/etc/modprobe.d/` directory

Example:

```
# printf '\n%s\n' "install atm /bin/false" >> /etc/modprobe.d/60-atm.conf
```

Create a file ending in `.conf` with `blacklist atm` in the `/etc/modprobe.d/` directory

Example:

```
# printf '\n%s\n' "blacklist atm" >> /etc/modprobe.d/60-atm.conf
```

References:

1. GRID: NIX-00000011
2. CCI-000381
3. NIST SP 800-53 :: CM-7
4. NIST SP 800-53A :: CM-7.1 (ii)
5. STIG ID: RHEL-08-040021 | RULE ID: SV-230494r1069310 | CAT III
6. STIG ID: OL08-00-040021 | RULE ID: SV-248829r991589 | CAT II
7. STIG ID: RHEL-09-213045 | RULE ID: SV-257804r1044853 | CAT II

Internal Only - General

8. STIG ID: ALMA-09-029610 | RULE ID: SV-269341r1050223 | CAT II
9. STIG ID: OL09-00-000040 | RULE ID: SV-271443r1092463 | CAT II

CIS Controls:

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

Internal Only - General

3.2.2 Ensure can kernel module is not available (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The Controller Area Network (CAN) is a serial communications protocol, which was initially developed for automotive and is now also used in marine, industrial, and medical applications. Disabling CAN protects the system against exploitation of any flaws in its implementation.

Rationale:

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

Audit:

Verify the `can` kernel module is not available on the system or has been disabled.

1. Run the following script to determine if the `can` kernel module is available on the system:

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

{
   l_mod_name="can" l_mod_type="net"
   while IFS= read -r l_mod_path; do
      if [ -d "$l_mod_path/${l_mod_name//-/\/}" ] &&  \
      [ -n "$(ls -A "$l_mod_path/${l_mod_name//-/\/}")" ]; then
         printf '%s\n' "$l_mod_name exists in $l_mod_path"
      fi
   done < <(readlink -e /usr/lib/modules/**/kernel/$l_mod_type \
   || readlink -e /lib/modules/**/kernel/$l_mod_type)
}
```

If nothing is returned, the `can` kernel module is not available on the system and no further audit steps are required.

Note: Some systems may include the `can` filesystem as part of the kernel opposed to being available as a kernel module. In this case, the above audit will not return anything. This is also considered a passing state.

If anything is returned by the above script:

2. verify the `can` kernel module is not loaded and not loadable by performing the following:

Internal Only - General

Run the following command to verify the `can` kernel module is not loaded:

```
# lsmod | grep 'can'
```

Nothing should be returned

Run the following command to verify the `can` kernel module is not loadable:

```
# modprobe --showconfig | grep -P -- '\b(install|blacklist)\h+can\b'
```

Verify the output includes:

```
blacklist can
 -AND-
install can /bin/false
 -OR-
install can /bin/true
```

Example output:

```
blacklist can
install can /bin/false
```

Remediation:

Unload and disable the `can` kernel module.

1. Run the following commands to unload the `can` kernel module:

```
# modprobe -r can 2>/dev/null
# rmmod can 2>/dev/null
```

2. Perform the following to disable the `can` kernel module:

Create a file ending in `.conf` with `install can /bin/false` in the `/etc/modprobe.d/` directory

Example:

```
# printf '\n%s\n' "install can /bin/false" >> /etc/modprobe.d/60-can.conf
```

Create a file ending in `.conf` with `blacklist can` in the `/etc/modprobe.d/` directory

Example:

```
# printf '\n%s\n' "blacklist can" >> /etc/modprobe.d/60-can.conf
```

References:

1. GRID: NIX-00000012
2. CCI-000381
3. NIST SP 800-53 :: CM-7
4. NIST SP 800-53A :: CM-7.1 (ii)
5. STIG ID: RHEL-08-040022 | RULE ID: SV-230495r1069311 | CAT III

Internal Only - General

6. STIG ID: OL08-00-040022 | RULE ID: SV-248830r991589 | CAT II
7. STIG ID: RHEL-09-213050 | RULE ID: SV-257805r1044856 | CAT II
8. STIG ID: ALMA-09-029830 | RULE ID: SV-269343r1050225 | CAT II
9. STIG ID: OL09-00-000041 | RULE ID: SV-271444r1091044 | CAT II

CIS Controls:

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

Internal Only - General

3.2.3 Ensure dccp kernel module is not available (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - 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:

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

Audit:

Verify the `dccp` kernel module is not available on the system or has been disabled.

1. Run the following script to determine if the `dccp` kernel module is available on the system:

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

{
   l_mod_name="dccp" l_mod_type="net"
   while IFS= read -r l_mod_path; do
      if [ -d "$l_mod_path/${l_mod_name//-/\/}" ] &&  \
      [ -n "$(ls -A "$l_mod_path/${l_mod_name//-/\/}")" ]; then
         printf '%s\n' "$l_mod_name exists in $l_mod_path"
      fi
   done < <(readlink -e /usr/lib/modules/**/kernel/$l_mod_type \
   || readlink -e /lib/modules/**/kernel/$l_mod_type)
}
```

If nothing is returned, the `dccp` kernel module is not available on the system and no further audit steps are required.

Note: Some systems may include the `dccp` filesystem as part of the kernel opposed to being available as a kernel module. In this case, the above audit will not return anything. This is also considered a passing state.

If anything is returned by the above script:

2. verify the `dccp` kernel module is not loaded and not loadable by performing the following:

Internal Only - General

Run the following command to verify the `dccp` kernel module is not loaded:

```
# lsmod | grep 'dccp'
```

Nothing should be returned

Run the following command to verify the `dccp` kernel module is not loadable:

```
# modprobe --showconfig | grep -P -- '\b(install|blacklist)\h+dccp\b'
```

Verify the output includes:

```
blacklist dccp
 -AND-
install dccp /bin/false
 -OR-
install dccp /bin/true
```

Example output:

```
blacklist dccp
install dccp /bin/false
```

Remediation:

Unload and disable the `dccp` kernel module.

1. Run the following commands to unload the `dccp` kernel module:

```
# modprobe -r dccp 2>/dev/null
# rmmod dccp 2>/dev/null
```

2. Perform the following to disable the `dccp` kernel module:

Create a file ending in `.conf` with `install dccp /bin/false` in the `/etc/modprobe.d/` directory

Example:

```
# printf '\n%s\n' "install dccp /bin/false" >> /etc/modprobe.d/60-dccp.conf
```

Create a file ending in `.conf` with `blacklist dccp` in the `/etc/modprobe.d/` directory

Example:

```
# printf '\n%s\n' "blacklist dccp" >> /etc/modprobe.d/60-dccp.conf
```

References:

1. GRID: NIX-00000013
2. NIST SP 800-53 :: CM-7
3. NIST SP 800-53A :: CM-7.1 (ii)
4. CCI-000381

Internal Only - General

CIS Controls:

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

Internal Only - General

3.2.4 Ensure rds kernel module is not available (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - 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:

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

Audit:

Verify the `rds` kernel module is not available on the system or has been disabled.

1. Run the following script to determine if the `rds` kernel module is available on the system:

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

{
   l_mod_name="rds" l_mod_type="net"
   while IFS= read -r l_mod_path; do
      if [ -d "$l_mod_path/${l_mod_name//-/\/}" ] &&  \
      [ -n "$(ls -A "$l_mod_path/${l_mod_name//-/\/}")" ]; then
         printf '%s\n' "$l_mod_name exists in $l_mod_path"
      fi
   done < <(readlink -e /usr/lib/modules/**/kernel/$l_mod_type \
   || readlink -e /lib/modules/**/kernel/$l_mod_type)
}
```

If nothing is returned, the `rds` kernel module is not available on the system and no further audit steps are required.

Note: Some systems may include the `rds` filesystem as part of the kernel opposed to being available as a kernel module. In this case, the above audit will not return anything. This is also considered a passing state.

If anything is returned by the above script:

2. verify the `rds` kernel module is not loaded and not loadable by performing the following:

Run the following command to verify the `rds` kernel module is not loaded:

Internal Only - General

```
# lsmod | grep 'rds'
```

Nothing should be returned

Run the following command to verify the `rds` kernel module is not loadable:

```
# modprobe --showconfig | grep -P -- '\b(install|blacklist)\h+rds\b'
```

Verify the output includes:

```
blacklist rds
 -AND-
install rds /bin/false
 -OR-
install rds /bin/true
```

Example output:

```
blacklist rds
install rds /bin/false
```

Remediation:

Unload and disable the `rds` kernel module.

1. Run the following commands to unload the `rds` kernel module:

```
# modprobe -r rds 2>/dev/null
# rmmod rds 2>/dev/null
```

2. Perform the following to disable the `rds` kernel module:

Create a file ending in `.conf` with `install rds /bin/false` in the `/etc/modprobe.d/` directory

Example:

```
# printf '\n%s\n' "install rds /bin/false" >> /etc/modprobe.d/60-rds.conf
```

Create a file ending in `.conf` with `blacklist rds` in the `/etc/modprobe.d/` directory

Example:

```
# printf '\n%s\n' "blacklist rds" >> /etc/modprobe.d/60-rds.conf
```

References:

1. GRID: NIX-00000014
2. NIST SP 800-53 :: CM-7
3. NIST SP 800-53A :: CM-7.1 (ii)
4. CCI-000381

Internal Only - General

CIS Controls:

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

Internal Only - General

3.2.5 Ensure sctp kernel module is not available (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - 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:

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

Audit:

Verify the `sctp` kernel module is not available on the system or has been disabled.

1. Run the following script to determine if the `sctp` kernel module is available on the system:

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

{
   l_mod_name="sctp" l_mod_type="net"
   while IFS= read -r l_mod_path; do
      if [ -d "$l_mod_path/${l_mod_name//-/\/}" ] &&  \
      [ -n "$(ls -A "$l_mod_path/${l_mod_name//-/\/}")" ]; then
         printf '%s\n' "$l_mod_name exists in $l_mod_path"
      fi
   done < <(readlink -e /usr/lib/modules/**/kernel/$l_mod_type \
   || readlink -e /lib/modules/**/kernel/$l_mod_type)
}
```

If nothing is returned, the `sctp` kernel module is not available on the system and no further audit steps are required.

Note: Some systems may include the `sctp` filesystem as part of the kernel opposed to being available as a kernel module. In this case, the above audit will not return anything. This is also considered a passing state.

If anything is returned by the above script:

2. verify the `sctp` kernel module is not loaded and not loadable by performing the following:

Internal Only - General

Run the following command to verify the `sctp` kernel module is not loaded:

```
# lsmod | grep 'sctp'
```

Nothing should be returned

Run the following command to verify the `sctp` kernel module is not loadable:

```
# modprobe --showconfig | grep -P -- '\b(install|blacklist)\h+sctp\b'
```

Verify the output includes:

```
blacklist sctp
 -AND-
install sctp /bin/false
 -OR-
install sctp /bin/true
```

Example output:

```
blacklist sctp
install sctp /bin/false
```

Remediation:

Unload and disable the `sctp` kernel module.

1. Run the following commands to unload the `sctp` kernel module:

```
# modprobe -r sctp 2>/dev/null
# rmmod sctp 2>/dev/null
```

2. Perform the following to disable the `sctp` kernel module:

Create a file ending in `.conf` with `install sctp /bin/false` in the `/etc/modprobe.d/` directory

Example:

```
# printf '\n%s\n' "install sctp /bin/false" >> /etc/modprobe.d/60-sctp.conf
```

Create a file ending in `.conf` with `blacklist sctp` in the `/etc/modprobe.d/` directory

Example:

```
# printf '\n%s\n' "blacklist sctp" >> /etc/modprobe.d/60-sctp.conf
```

References:

1. GRID: NIX-00000015
2. NIST SP 800-53 :: CM-7
3. NIST SP 800-53A :: CM-7.1 (ii)
4. STIG ID: RHEL-08-040023 | RULE ID: SV-230496r1069312 | CAT III
5. STIG ID: OL08-00-040023 | RULE ID: SV-248831r991589 | CAT II

Internal Only - General

6. STIG ID: RHEL-09-213060 | RULE ID: SV-257807r1044862 | CAT II
7. STIG ID: ALMA-09-030050 | RULE ID: SV-269345r1050227 | CAT II
8. STIG ID: OL09-00-000043 | RULE ID: SV-271446r1091050 | CAT II
9. CCI-000381

CIS Controls:

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

Internal Only - General

3.2.6 Ensure tipc kernel module is not available (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The Transparent Inter-Process Communication (TIPC) protocol is designed to provide communication between cluster nodes.

Rationale:

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

Audit:

Verify the `tipc` kernel module is not available on the system or has been disabled.

1. Run the following script to determine if the `tipc` kernel module is available on the system:

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

{
   l_mod_name="tipc" l_mod_type="net"
   while IFS= read -r l_mod_path; do
      if [ -d "$l_mod_path/${l_mod_name//-/\/}" ] &&  \
      [ -n "$(ls -A "$l_mod_path/${l_mod_name//-/\/}")" ]; then
         printf '%s\n' "$l_mod_name exists in $l_mod_path"
      fi
   done < <(readlink -e /usr/lib/modules/**/kernel/$l_mod_type \
   || readlink -e /lib/modules/**/kernel/$l_mod_type)
}
```

If nothing is returned, the `tipc` kernel module is not available on the system and no further audit steps are required.

Note: Some systems may include the `tipc` filesystem as part of the kernel opposed to being available as a kernel module. In this case, the above audit will not return anything. This is also considered a passing state.

If anything is returned by the above script:

2. verify the `tipc` kernel module is not loaded and not loadable by performing the following:

Run the following command to verify the `tipc` kernel module is not loaded:

Internal Only - General

```
# lsmod | grep 'tipc'
```

Nothing should be returned

Run the following command to verify the `tipc` kernel module is not loadable:

```
# modprobe --showconfig | grep -P -- '\b(install|blacklist)\h+tipc\b'
```

Verify the output includes:

```
blacklist tipc
 -AND-
install tipc /bin/false
 -OR-
install tipc /bin/true
```

Example output:

```
blacklist tipc
install tipc /bin/false
```

Remediation:

Unload and disable the `tipc` kernel module.

1. Run the following commands to unload the `tipc` kernel module:

```
# modprobe -r tipc 2>/dev/null
# rmmod tipc 2>/dev/null
```

2. Perform the following to disable the `tipc` kernel module:

Create a file ending in `.conf` with `install tipc /bin/false` in the `/etc/modprobe.d/` directory

Example:

```
# printf '\n%s\n' "install tipc /bin/false" >> /etc/modprobe.d/60-tipc.conf
```

Create a file ending in `.conf` with `blacklist tipc` in the `/etc/modprobe.d/` directory

Example:

```
# printf '\n%s\n' "blacklist tipc" >> /etc/modprobe.d/60-tipc.conf
```

References:

1. GRID: NIX-00000016
2. NIST SP 800-53 :: CM-7
3. NIST SP 800-53A :: CM-7.1 (ii)
4. STIG ID: RHEL-08-040024 | RULE ID: SV-230497r1069313 | CAT III
5. STIG ID: OL08-00-040024 | RULE ID: SV-248832r958478 | CAT III
6. STIG ID: RHEL-09-213065 | RULE ID: SV-257808r1044865 | CAT II
7. STIG ID: ALMA-09-030270 | RULE ID: SV-269347r1050229 | CAT II

Internal Only - General

8. STIG ID: OL09-00-000044 | RULE ID: SV-271447r1092464 | CAT II
9. CCI-000381

CIS Controls:

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

Internal Only - General

3.2.7 Ensure unneeded network protocol kernel modules are not available (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The Linux kernel ships dozens of network-protocol implementations as loadable modules under `kernel/net/` (e.g., `esp4`, `esp6`, `rxrpc`, `sctp`, `dccp`, `rds`, `tipc`, `appletalk`, `decnet`, `ipx`, `x25`, `rose`, `ax25`, `netrom`, `econet`, `llc`, `llc2`, and many others). On a default installation most of these modules are **available but not loaded** — present on disk under `/lib/modules/$(uname -r)/kernel/net/`, autoloadable on demand by `modprobe`, by socket-family registration (`AF_*` socket calls), or by traffic arriving on an interface that requests the protocol.

This recommendation establishes a **site-maintained allow-list / deny-list policy** for network-protocol kernel modules: the operator enumerates every network-protocol module the site does **not** need and ensures each is blacklisted and bound to `/bin/false` (or `/bin/true`) via `/etc/modprobe.d/`, so the kernel will refuse to autoload or manually load them.

This rule is the **catch-all framework** that complements the per-protocol recommendations elsewhere in subsection 3.2 Configure Network Kernel Modules.

1. The set of "unneeded" network-protocol modules differs materially between deployments (e.g., a site using IPsec needs `esp4`/`esp6` loaded; an AFS-using site needs `rxrpc`; most sites need neither).
2. New network-protocol modules and new vulnerabilities in existing modules are disclosed continuously, and a hard-coded list in a CIS recommendation cannot keep pace with the kernel's release cadence or with public vulnerability disclosure.
3. A site that has already inventoried its required network protocols can apply this framework immediately when a new vulnerability is published, without waiting for a benchmark revision.

Internal Only - General

Worked example (illustrative, not normative): On May 7, 2026 Canonical disclosed the "Dirty Frag" local privilege escalation class — CVE-2026-43284 in the `esp4`/`esp6` modules (IPsec ESP fragmentation handling) and CVE-2026-43500 in the `rxrpc` module (Andrew File System RPC) — affecting all Ubuntu LTS releases from 14.04 through 26.04. A site that did not use IPsec or AFS could mitigate both CVEs immediately by adding `esp4`, `esp6`, and `rxrpc` to its `/etc/modprobe.d/` deny list — exactly the workflow this recommendation institutionalises. A site that did use IPsec or AFS would scope those modules out of its deny list and accept the standard patching cycle for them instead.

This recommendation does not prescribe a specific list of modules. It prescribes the **discipline** of maintaining and enforcing one.

Rationale:

Every loadable network-protocol kernel module is a piece of code that can be reached by an unprivileged local user (via `socket(AF_FAMILY, ...)` for protocols mapped to a socket family, via `ip link add type <kind>` for some encapsulation modules, or via autoload triggers from received traffic). The protocol code typically runs in the kernel at the highest privilege level. A memory-safety bug, integer-overflow, or fragmentation-reassembly defect anywhere in this code is exploitable for local privilege escalation or, in some cases, remote denial-of-service or remote code execution.

The 2026 "Dirty Frag" disclosure is one example of this class — local unprivileged users on Ubuntu LTS systems could elevate to root by exercising vulnerable ESP fragmentation handling in `esp4`/`esp6` (CVE-2026-43284, CVSS 8.8) or AFS RPC processing in `rxrpc` (CVE-2026-43500, CVSS 7.8). Container escapes were also plausible because both modules expose attack surface to unprivileged user namespaces. Similar disclosures have historically affected `dccp`, `sctp`, `rds`, `tipc`, `vsock`, `kcm`, `phonet`, and others; the cadence is sustained.

Sites that have not pre-emptively disabled modules they don't use are forced into reactive patching every time such a disclosure lands — accepting a window of exposure between disclosure and patch-deployment. Sites that maintain an explicit, enforced deny list under `/etc/modprobe.d/` close the attack surface in advance, so any newly disclosed vulnerability in a module they have already disabled has zero impact on them.

This is the same defence-in-depth posture taken by recommendations 3.2.1 through 3.2.6, generalised to give site administrators a documented, auditable workflow for protocols not individually called out by the benchmark.

Impact:

Disabling a network-protocol kernel module prevents any process or kernel subsystem from using that protocol. Before adding a module to the site's deny list, the operator must verify that no production workload depends on it. Common dependency examples:

- `esp4` / `esp6` — required for IPsec/IKE VPNs (strongSwan, Libreswan, kernel `xfrm`). Disabling breaks site-to-site VPNs and `ipsec`-mode WireGuard interop.

Internal Only - General

- `rxrpc` — required for Andrew File System (`afs`). Disabling breaks AFS clients and any `kafs` mounts.
- `mptcp` — required for Multipath TCP. Disabling reverts MPTCP connections to plain TCP single-path.
- `tcp_bbr` / `tcp_cubic` — congestion-control modules; disabling the active one falls back to compiled-in defaults but may degrade throughput on long-fat-network paths.
- `nf_conntrack_*` / `nf_nat_*` — netfilter helpers for specific protocols (`ftp`, `sip`, `pptp`, etc.); disabling breaks NAT traversal for those protocols.

Operators must inventory required network protocols before applying the deny list. A reboot is recommended after deny-list changes to confirm no unit fails to start because its required protocol module is now blocked; alternatively, `update-initramfs -u` followed by service-by-service restart can validate the configuration without rebooting.

The deny-list file under `/etc/modprobe.d/` survives package upgrades but is **not** automatically populated as new vulnerable modules are disclosed — the operator is responsible for reviewing the site's allow/deny inventory on the same cadence as the site's vulnerability-management process.

Audit:

This recommendation is **Manual**. The site must:

1. Maintain a documented list of network-protocol kernel modules that are **not** required by any production workload. The list lives in the site's configuration-management baseline and is reviewed on the same cadence as vulnerability disclosures.
2. Enforce that list via one or more files in `/etc/modprobe.d/` ending in `.conf` (conventionally `60-cis-network-disabled.conf` or split per-protocol per CIS examples).
3. Verify the enforcement matches the site list.

For **each** network-protocol kernel module the site has determined is unneeded, the auditor must confirm the module is unavailable on the system. Substituting `<module>` with the module name from the site list:

Run the following command to verify the `<module>` kernel module is not loaded:

```
# lsmod | grep -E '^<module>\s'
```

Nothing should be returned.

Run the following command to verify the `<module>` kernel module is bound to `/bin/false` or `/bin/true` in the effective modprobe configuration:

```
# modprobe --showconfig | grep -P --
'^\s*install\s+<module>\s+/bin/(false|true)\b'
```

Verify the output is one of:

Internal Only - General

```
install <module> /bin/false
install <module> /bin/true
```

Run the following command to verify the `<module>` kernel module is blacklisted:

```
# modprobe --showconfig | grep -P -- '^\s*blacklist\s+<module>\b'
```

Verify the output is:

```
blacklist <module>
```

Run the following command to confirm a manual load attempt is refused:

```
# modprobe -n -v <module> 2>&1
```

Verify the output contains `install /bin/false` (or `install /bin/true`) and does **not** contain `insmod` against a real `.ko` file.

Note: This audit must be repeated for every entry on the site's "unneeded network protocols" list. The per-protocol audits required by recommendations 3.2.1 through 3.2.6 satisfy the framework of this rule for those six modules; this rule covers every **additional** module the site has placed on its deny list.

Note: If the site has formally documented an empty deny list — i.e., it has reviewed every network-protocol module shipped by the running kernel and determined every one is required — this recommendation is satisfied by producing that signed inventory. In practice this is rare; most sites will have at least a handful of clearly unneeded protocols.

Remediation:

1. Produce a site inventory of required network-protocol kernel modules. List each `kernel/net/*` module currently shipped with the running kernel and tag each as **required**, **unneeded**, or **already covered by 3.2.1–3.2.6**.

To enumerate the modules shipped with the running kernel:

```
# find /lib/modules/"$(uname -r)"/kernel/net -type f -name '*.ko*' -printf
'%f\n' | sed -E 's/\.ko(\.[xg]z|\.zst)?$//' | sort -u
```

2. For each module flagged **unneeded**, append both an `install` and a `blacklist` directive to a single site-managed file under `/etc/modprobe.d/`. The conventional filename used by the per-protocol recs in this section is `60-<module>.conf`; the catch-all may be consolidated into a single file (e.g., `60-cis-network-disabled.conf`) or split per-module.

Example, consolidating multiple modules into one file:

Internal Only - General

```
# cat > /etc/modprobe.d/60-cis-network-disabled.conf << 'EOF'
# CIS 3.2.7 — site-maintained network-protocol module deny list
# Each entry must be reviewed against the site's required-protocol inventory.

install esp4 /bin/false
blacklist esp4

install esp6 /bin/false
blacklist esp6

install rxrpc /bin/false
blacklist rxrpc

# Add additional unneeded network-protocol modules here.
EOF
```

3. Unload any of the listed modules that are currently loaded:

```
# for m in esp4 esp6 rxrpc; do modprobe -r "$m" 2>/dev/null; rmmod "$m"
2>/dev/null; done
```

4. Regenerate the initramfs so the deny list applies at boot, before any subsystem auto-loads a denied module:

```
# update-initramfs -u
```

5. Confirm the configuration is in effect by re-running the Audit Procedure against each module on the site list.
6. Record the site list in the configuration-management baseline (Ansible, Puppet, Chef, Salt, or equivalent) so the deny-list file is reproduced on rebuild and so additions are reviewed via the site's change-control process.

Note: `install <module> /bin/false` is preferred over `blacklist` alone. `blacklist` only prevents *autoload* triggered by alias resolution; an explicit `modprobe <module>` still succeeds. Pairing `blacklist` with `install <module> /bin/false` blocks both autoload and manual load. The per-protocol recs 3.2.1–3.2.6 use the same pattern.

References:

1. GRID: NIX-00000407
2. NIST SP 800-53 Rev. 5: CM-7, CM-6
3. https://manpages.debian.org/trixie/kmod/modprobe.d.5.en.html
4. https://manpages.debian.org/trixie/kmod/modprobe.8.en.html
5. https://canonical.com/blog/dirty-frag-linux-vulnerability-fixes-available
6. CVE-2026-43284 (Dirty Frag — ESP, IPsec fragmentation handling)
7. CVE-2026-43500 (Dirty Frag — RxRPC, AFS RPC handling)

Internal Only - General

CIS Controls:

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

3.3 Configure Network Kernel Parameters

Note: Some of the Recommendations in this section's subsections are not applicable if the system is a router. These Recommendations include a note in their description.

This section contains Recommendations covering kernel parameters affecting the network stack.

- sysctl settings are defined through files in `/usr/local/lib`, `/usr/lib/`, `/lib/`, `/run/`, and `/etc/`
- Files are typically placed in the `sysctl.d` directory within the parent directory
- The paths where sysctl preload files usually exist
  - `/run/sysctl.d/*.conf`
  - `/etc/sysctl.d/*.conf`
  - `/usr/local/lib/sysctl.d/*.conf`
  - `/usr/lib/sysctl.d/*.conf`
  - `/lib/sysctl.d/*.conf`
  - `/etc/sysctl.conf`
- Files must have the "`.conf`" extension
- Vendors settings usually live in `/usr/lib/` or `/usr/local/lib/`
- To override a whole file, create a new file with the same name in `/etc/sysctl.d/` and put new settings there.
- To override only specific settings, add a file with a lexically later name in `/etc/sysctl.d/` and put new settings there.
- The command `/usr/lib/systemd/systemd-sysctl --cat-config` produces output containing The system's loaded kernel parameters and the files they're configured in:
  - Entries listed latter in the file take precedence over the same settings listed earlier in the file

Internal Only - General

    - Files containing kernel parameters that are over-ridden by other files with the same name will not be listed
    - On systems running UncomplicatedFirewall, the kernel parameters may be set or over-written. This will not be visible in the output of the command
- On systems with Uncomplicated Firewall, additional settings may be configured in `/etc/ufw/sysctl.conf`
- The settings in `/etc/ufw/sysctl.conf` will override settings other settings and **will not** be visible in the output of the `/usr/lib/systemd/systemd-sysctl --cat-config` command
- This behavior can be changed by updating the `IPT_SYSCTL` parameter in `/etc/default/ufw`
- **SYSTEM FILE PRECEDENCE**
    - When using the `--system` option, `sysctl` will read files from directories in the following list in given order from top to bottom. Once a file of a given filename is loaded, any file of the same name in subsequent directories is ignored.

      ```
      /etc/sysctl.d/*.conf /run/sysctl.d/*.conf
      /usr/local/lib/sysctl.d/*.conf /usr/lib/sysctl.d/*.conf
      /lib/sysctl.d/*.conf /etc/sysctl.conf
      ```

    - All configuration files are sorted in lexicographic order, regardless of the directory they reside in. Configuration files can either be completely replaced (by having a new configuration file with the same name in a directory of higher priority) or partially replaced (by having a configuration file that is ordered later)--

Internal Only - General

3.3.1 Configure IPv4 parameters

The recommendations in this subsection configure `net.ipv4` kernel parameters that govern the behavior of the IPv4 network stack on Debian Linux systems. These parameters control packet forwarding, ICMP handling, redirect processing, reverse path filtering, source route acceptance, and logging of anomalous packets. Each parameter is applied via the `sysctl` facility and persists across reboots when set in a file under `/etc/sysctl.d/` or `/etc/sysctl.conf`.

More information about the kernel parameter configuration files, their location, and load order is available in the **Configure Network Kernel Parameters** section overview.

Router applicability

If this system is configured as a network router, the recommendations in this subsection are not applicable.

Parameter interdependencies

Several parameters in this subsection are interdependent. When implementing, note the following:

- If `net.ipv4.ip_forward` is set to `0` (3.3.1.1), recommendations 3.3.1.2 (`net.ipv4.conf.all.forwarding`) and 3.3.1.3 (`net.ipv4.conf.default.forwarding`) may be skipped, as the forwarding behavior is effectively superseded.
- Conversely, if both `net.ipv4.conf.all.forwarding=0` and `net.ipv4.conf.default.forwarding=0` are confirmed, 3.3.1.1 (`net.ipv4.ip_forward`) may be skipped.

UFW sysctl file interaction

Systems with Uncomplicated Firewall (UFW) installed maintain a separate sysctl configuration file, referenced by the `IPT_SYSCTL` variable in `/etc/default/ufw`. The audit and remediation scripts in this subsection include this file in their scan of sysctl configuration files and will comment out any non-compliant parameter values found there.

However, subsequent UFW package updates may rewrite this file and re-enable previously corrected values. When that occurs, the UFW sysctl file entry takes precedence over the value written to `/etc/sysctl.d/` because UFW's file is loaded earlier in the systemd sysctl configuration priority order. After any UFW package update, verify that the parameters in this subsection retain their expected values by re-running the audit procedures.

Internal Only - General

3.3.1.1 Ensure net.ipv4.ip_forward is configured (Automated)

Profile Applicability:

- Level 1 - Workstation
- Level 1 - Server

Description:

The `net.ipv4.ip_forward` flag tells the system whether it can forward IPv4 packets or not.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

Rationale:

Routing protocol daemons are typically used on routers to exchange network topology information with other routers. If this software is used when not required, system network information may be unnecessarily transmitted across the network.

Setting `net.ipv4.ip_forward` to `0` ensures that a system with multiple interfaces (for example, a hard proxy), will not be able to forward IPv4 packets.

Impact:

IP forwarding is required on systems configured to act as a router. If these parameters are disabled, the system will not be able to perform as a router.

Many Cloud Service Provider (CSP) hosted systems require IP forwarding to be enabled. If the system is running on a CSP platform, this requirement should be reviewed before disabling IP forwarding.

Audit:

Verify `net.ipv4.ip_forward` is set to `0`.

1. Run the following command to verify `net.ipv4.ip_forward` is set to `0` in the running configuration:

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

Verify output is:

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

2. Run the following script to verify `net.ipv4.ip_forward = 0` is set in a file being used by systemd sysctl to configure `net.ipv4.ip_forward`:

Internal Only - General

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

{
   l_parameter_name="net.ipv4.ip_forward"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv4.ip_forward = 0`.

Example output:

```
 - "net.ipv4.ip_forward = 0" is set in: "/etc/sysctl.d/60-ipv4_sysctl.conf"
```

Note: If multiple lines with `net.ipv4.ip_forward =` are returned.

- The first line includes the value being used by `systemd-sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect value to minimize the potential of the incorrect value being used by `systemd-sysctl` due to system configuration changes.

Remediation:

1. Review all files being used by systemd-sysctl and comment out or remove all `net.ipv4.ip_forward` lines that are not `net.ipv4.ip_forward=0`.

Example script:

Internal Only - General

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

{
   l_option="net.ipv4.ip_forward" l_value="0"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

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

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv4.ip_forward = 0" >> /etc/sysctl.d/60-
ipv4_sysctl.conf
```

Note: See the Configure IPv4 parameters subsection overview for important notes on UFW sysctl file interaction.

3. Run the following command to load all sysctl configuration files:

```
# sysctl --system
```

Default Value:

net.ipv4.ip_forward = 1

References:

1. GRID: NIX-00000017
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: SLES-15-040380 | RULE ID: SV-235024r991589 | CAT II
5. CCI-000366

Internal Only - General

CIS Controls:

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

Internal Only - General

3.3.1.2 Ensure net.ipv4.conf.all.forwarding is configured (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.

`net.ipv4.conf.all.forwarding` controls forwarding of IPv4 packet on all interfaces.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

Rationale:

Routing protocol daemons are typically used on routers to exchange network topology information with other routers. If this capability is used when not required, system network information may be unnecessarily transmitted across the network.

Impact:

IP forwarding is required on systems configured to act as a router. If these parameters are disabled, the system will not be able to perform as a router.

Cloud Service Provider (CSP) hosted systems may require forwarding to be enabled. If the system is running on a CSP platform, this requirement should be reviewed before disabling IPv4 forwarding.

Audit:

Verify `net.ipv4.conf.all.forwarding` is set to `0`.

1. Run the following command to verify `net.ipv4.conf.all.forwarding` is set to `0` in the running configuration:

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

Verify output is:

```
net.ipv4.conf.all.forwarding = 0
```

2. Run the following script to verify `net.ipv4.conf.all.forwarding = 0` is set in a file being used by systemd-sysctl to configure `net.ipv4.conf.all.forwarding`:

Internal Only - General

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

{
   l_parameter_name="net.ipv4.conf.all.forwarding"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv4.conf.all.forwarding = 0`.

Example output:

```
 - "net.ipv4.conf.all.forwarding = 0" is set in: "/etc/sysctl.d/60-
ipv4_sysctl.conf"
```

Note:

- If multiple lines with `net.ipv4.conf.all.forwarding =` are returned.
- The first line includes the value being used by `systemd-sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect value to minimize the potential of the incorrect value being used by `systemd-sysctl` due to system configuration changes.
- This script looks at all files used by systemd sysctl.
- More information about these files, their location, and load preference is available in the section overview.

Internal Only - General

Remediation:

1. Review all files being used by systemd-sysctl and comment out or remove all `net.ipv4.conf.all.forwarding` lines that are not `net.ipv4.conf.all.forwarding=0`.

Example script:

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

{
   l_option="net.ipv4.conf.all.forwarding" l_value="0"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

```
net.ipv4.conf.all.forwarding = 0
```

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv4.conf.all.forwarding = 0" \
>> /etc/sysctl.d/60-ipv4_sysctl.conf
```

Note: See the Configure IPv4 parameters subsection overview for important notes on UFW sysctl file interaction.

3. Run the following command to load all sysctl configuration files:

```
# sysctl --system
```

Default Value:

net.ipv4.conf.all.forwarding = 1

Internal Only - General

References:

1. GRID: NIX-00000018
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-08-040259 | RULE ID: SV-250317r1017358 | CAT II
5. STIG ID: OL08-00-040259 | RULE ID: SV-252662r991589 | CAT II
6. STIG ID: RHEL-09-253075 | RULE ID: SV-257970r1045011 | CAT II
7. STIG ID: ALMA-09-019160 | RULE ID: SV-269249r1050131 | CAT II
8. STIG ID: OL09-00-006028 | RULE ID: SV-271872r1092328 | CAT II
9. CCI-000366

CIS Controls:

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

Internal Only - General

3.3.1.3 Ensure net.ipv4.conf.default.forwarding is configured (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.

`net.ipv4.conf.default.forwarding` controls forwarding of IPv4 packet on newly added network interfaces.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

Rationale:

Routing protocol daemons are typically used on routers to exchange network topology information with other routers. If this capability is used when not required, system network information may be unnecessarily transmitted across the network.

Impact:

IP forwarding is required on systems configured to act as a router. If these parameters are disabled, the system will not be able to perform as a router.

Cloud Service Provider (CSP) hosted systems may require forwarding to be enabled. If the system is running on a CSP platform, this requirement should be reviewed before disabling IPv4 forwarding.

Audit:

Verify `net.ipv4.conf.default.forwarding` is set to `0`.

1. Run the following command to verify `net.ipv4.conf.default.forwarding` is set to `0` in the running configuration:

```
# sysctl net.ipv4.conf.default.forwarding
```

Verify output is:

Internal Only - General

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

2. Run the following script to verify `net.ipv4.conf.default.forwarding = 0` is set in a file being used by systemd sysctl to configure `net.ipv4.conf.default.forwarding`:

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

{
   l_parameter_name="net.ipv4.conf.default.forwarding"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv4.conf.default.forwarding = 0`.

Example output:

```
 - "net.ipv4.conf.default.forwarding = 0" is set in: "/etc/sysctl.d/60-
ipv4_sysctl.conf"
```

Note: If multiple lines with `net.ipv4.conf.default.forwarding =` are returned.

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect

Internal Only - General

value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `net.ipv4.conf.default.forwarding` lines that are not `net.ipv4.conf.default.forwarding=0`.

Example script:

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

{
   l_option="net.ipv4.conf.default.forwarding" l_value="0"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

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

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv4.conf.default.forwarding = 0" \
>> /etc/sysctl.d/60-ipv4_sysctl.conf
```

Note: See the Configure IPv4 parameters subsection overview for important notes on UFW sysctl file interaction.

3. Run the following command to load all sysctl configuration filles:

Internal Only - General

```
# sysctl --system
```

Default Value:

net.ipv4.conf.default.forwarding = 1

References:

1. GRID: NIX-00000019
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: ALMA-09-019160 | RULE ID: SV-269249r1050131 | CAT II
5. CCI-000366

CIS Controls:

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

Internal Only - General

3.3.1.4 Ensure net.ipv4.conf.all.send_redirects is configured (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.

`net.ipv4.conf.all.send_redirects` controls sending of all IPv4 ICMP redirected packets on all interfaces.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

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.

Impact:

IP forwarding is required on systems configured to act as a router. If these parameters are disabled, the system will not be able to perform as a router.

Audit:

Verify `net.ipv4.conf.all.send_redirects` is set to `0`.

1. Run the following command to verify `net.ipv4.conf.all.send_redirects` is set to `0` in the running configuration:

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

Verify output is:

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

2. Run the following script to verify `net.ipv4.conf.all.send_redirects = 0` is set in a file being used by systemd sysctl to configure `net.ipv4.conf.all.send_redirects`:

Internal Only - General

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

{
   l_parameter_name="net.ipv4.conf.all.send_redirects"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv4.conf.all.send_redirects = 0`.

Example output:

```
 - "net.ipv4.conf.all.send_redirects = 0" is set in: "/etc/sysctl.d/60-
ipv4_sysctl.conf"
```

Note: If multiple lines with `net.ipv4.conf.all.send_redirects =` are returned.

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Internal Only - General

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `net.ipv4.conf.all.send_redirects` lines that are not `net.ipv4.conf.all.send_redirects=0`.

Example script:

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

{
   l_option="net.ipv4.conf.all.send_redirects" l_value="0"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

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

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv4.conf.all.send_redirects = 0" \
>> /etc/sysctl.d/60-ipv4_sysctl.conf
```

Note: See the Configure IPv4 parameters subsection overview for important notes on UFW sysctl file interaction.

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv4.conf.all.send_redirects = 1

Internal Only - General

References:

1. GRID: NIX-00000020
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-08-040220 | RULE ID: SV-230536r1017298 | CAT II
5. STIG ID: OL08-00-040220 | RULE ID: SV-248877r991589 | CAT II
6. STIG ID: RHEL-09-253065 | RULE ID: SV-257968r991589 | CAT II
7. STIG ID: ALMA-09-019930 | RULE ID: SV-269256r1050138 | CAT II
8. STIG ID: OL09-00-006032 | RULE ID: SV-271875r1092337 | CAT II
9. STIG ID: SLES-15-040370 | RULE ID: SV-235023r991589 | CAT II
10. CCI-000366

CIS Controls:

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

Internal Only - General

3.3.1.5 Ensure net.ipv4.conf.default.send_redirects is configured (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.

`net.ipv4.conf.default.send_redirects` controls sending of all IPv4 ICMP redirected packets on a newly added network interface.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

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.

Impact:

IP forwarding is required on systems configured to act as a router. If these parameters are disabled, the system will not be able to perform as a router.

Audit:

Verify `net.ipv4.conf.default.send_redirects` is set to `0`.

1. Run the following command to verify `net.ipv4.conf.default.send_redirects` is set to `0` in the running configuration:

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

Verify output is:

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

2. Run the following script to verify `net.ipv4.conf.default.send_redirects = 0` is set in a file being used by systemd sysctl to configure `net.ipv4.conf.default.send_redirects`:

Internal Only - General

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

{
   l_parameter_name="net.ipv4.conf.default.send_redirects"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv4.conf.default.send_redirects = 0`.

Example output:

```
 - "net.ipv4.conf.default.send_redirects = 0" is set in: "/etc/sysctl.d/60-
ipv4_sysctl.conf"
```

Note: If multiple lines with `net.ipv4.conf.default.send_redirects =` are returned.

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Internal Only - General

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `net.ipv4.conf.default.send_redirects` lines that are not `net.ipv4.conf.default.send_redirects=0`.

Example script:

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

{
   l_option="net.ipv4.conf.default.send_redirects" l_value="0"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

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

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv4.conf.default.send_redirects = 0" \
>> /etc/sysctl.d/60-ipv4_sysctl.conf
```

Note: See the Configure IPv4 parameters subsection overview for important notes on UFW sysctl file interaction.

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv4.conf.default.send_redirects = 1

Internal Only - General

References:

1. GRID: NIX-00000021
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-08-040270 | RULE ID: SV-230543r1017305 | CAT II
5. STIG ID: OL08-00-040270 | RULE ID: SV-248886r991589 | CAT II
6. STIG ID: RHEL-09-253070 | RULE ID: SV-257969r991589 | CAT II
7. STIG ID: ALMA-09-019930 | RULE ID: SV-269256r1050138 | CAT II
8. STIG ID: OL09-00-006033 | RULE ID: SV-271876r1092641 | CAT II
9. STIG ID: SLES-15-040360 | RULE ID: SV-235022r991589 | CAT II
10. CCI-000366

CIS Controls:

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

Internal Only - General

3.3.1.6 Ensure net.ipv4.icmp_ignore_bogus_error_responses is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`net.ipv4.icmp_ignore_bogus_error_responses` controls if the kernel logs bogus responses (RFC-1122 non-compliant) from broadcast reframes.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

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:

Verify `net.ipv4.icmp_ignore_bogus_error_responses` is set to `1`.

1. Run the following command to verify `net.ipv4.icmp_ignore_bogus_error_responses` is set to `1` in the running configuration:

```
# sysctl net.ipv4.icmp_ignore_bogus_error_responses
```

Verify output is:

```
net.ipv4.icmp_ignore_bogus_error_responses = 1
```

2. Run the following script to verify `net.ipv4.icmp_ignore_bogus_error_responses = 1` is set in a file being used by systemd sysctl to configure `net.ipv4.icmp_ignore_bogus_error_responses`:

Internal Only - General

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

{
   l_parameter_name="net.ipv4.icmp_ignore_bogus_error_responses"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv4.icmp_ignore_bogus_error_responses = 1`.

Example output:

```
 - "net.ipv4.icmp_ignore_bogus_error_responses = 1" is set in:
"/etc/sysctl.d/60-ipv4_sysctl.conf"
```

Note: If multiple lines with `net.ipv4.icmp_ignore_bogus_error_responses =` are returned.

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Internal Only - General

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `net.ipv4.icmp_ignore_bogus_error_responses` lines that are not `net.ipv4.icmp_ignore_bogus_error_responses=1`.

Example script:

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

{
   l_option="net.ipv4.icmp_ignore_bogus_error_responses" l_value="1"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

```
net.ipv4.icmp_ignore_bogus_error_responses = 1
```

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv4.icmp_ignore_bogus_error_responses = 1" \
>> /etc/sysctl.d/60-ipv4_sysctl.conf
```

Note: See the Configure IPv4 parameters subsection overview for important notes on UFW sysctl file interaction.

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv4.icmp_ignore_bogus_error_responses = 1

Internal Only - General

References:

1. GRID: NIX-00000022
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-09-253060 | RULE ID: SV-257967r991589 | CAT II
5. STIG ID: ALMA-09-018940 | RULE ID: SV-269247r1050129 | CAT II
6. STIG ID: OL09-00-006031 | RULE ID: SV-271874r1092612 | CAT II
7. CCI-000366

CIS Controls:

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

Internal Only - General

3.3.1.7 Ensure net.ipv4.icmp_echo_ignore_broadcasts is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`net.ipv4.icmp_echo_ignore_broadcasts` controls if all ICMP echo and timestamp requests to broadcast and multicast addresses are ignored.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

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:

Verify `net.ipv4.icmp_echo_ignore_broadcasts` is set to `1`.

1. Run the following command to verify `net.ipv4.icmp_echo_ignore_broadcasts` is set to `1` in the running configuration:

```
# sysctl net.ipv4.icmp_echo_ignore_broadcasts
```

Verify output is:

```
net.ipv4.icmp_echo_ignore_broadcasts = 1
```

2. Run the following script to verify `net.ipv4.icmp_echo_ignore_broadcasts = 1` is set in a file being used by systemd sysctl to configure `net.ipv4.icmp_echo_ignore_broadcasts`:

Internal Only - General

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

{
   l_parameter_name="net.ipv4.icmp_echo_ignore_broadcasts"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv4.icmp_echo_ignore_broadcasts = 1`.

Example output:

```
 - "net.ipv4.icmp_echo_ignore_broadcasts = 1" is set in: "/etc/sysctl.d/60-
ipv4_sysctl.conf"
```

Note: If multiple lines with `net.ipv4.icmp_echo_ignore_broadcasts =` are returned.

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Internal Only - General

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `net.ipv4.icmp_echo_ignore_broadcasts` lines that are not `net.ipv4.icmp_echo_ignore_broadcasts=1`.

Example script:

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

{
   l_option="net.ipv4.icmp_echo_ignore_broadcasts" l_value="1"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

```
net.ipv4.icmp_echo_ignore_broadcasts = 1
```

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv4.icmp_echo_ignore_broadcasts = 1" \
>> /etc/sysctl.d/60-ipv4_sysctl.conf
```

Note: See the Configure IPv4 parameters subsection overview for important notes on UFW sysctl file interaction.

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv4.icmp_echo_ignore_broadcasts = 1

Internal Only - General

References:

1. GRID: NIX-00000023
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-08-040230 | RULE ID: SV-230537r1017299 | CAT II
5. STIG ID: OL08-00-040230 | RULE ID: SV-248878r991589 | CAT II
6. STIG ID: RHEL-09-253055 | RULE ID: SV-257966r991589 | CAT II
7. STIG ID: ALMA-09-019050 | RULE ID: SV-269248r1050130 | CAT II
8. STIG ID: OL09-00-006030 | RULE ID: SV-271873r1092331 | CAT II
9. CCI-000366

CIS Controls:

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

Internal Only - General

3.3.1.8 Ensure net.ipv4.conf.all.accept_redirects is configured (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.

`net.ipv4.conf.all.accept_redirects` controls accepting of all IPv4 ICMP redirected packets on all interfaces.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

Rationale:

ICMP redirect messages are packets that convey routing information and tell your host (acting as a router) to send packets via an alternate path. It is a way of allowing an outside routing device to update your system routing tables.

By setting `net.ipv4.conf.all.accept_redirects` to `0`, the system will not accept any IPv4 ICMP redirect messages, and therefore, won't allow outsiders to update the system's routing tables.

Audit:

Verify `net.ipv4.conf.all.accept_redirects` is set to `0`.

1. Run the following command to verify `net.ipv4.conf.all.accept_redirects` is set to `0` in the running configuration:

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

Verify output is:

```
net.ipv4.conf.all.accept_redirects = 0
```

2. Run the following script to verify `net.ipv4.conf.all.send_redirects = 0` is set in a file being used by systemd sysctl to configure `net.ipv4.conf.all.send_redirects`:

Internal Only - General

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

{
   l_parameter_name="net.ipv4.conf.all.accept_redirects"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv4.conf.all.accept_redirects = 0`.

Example output:

```
 - "net.ipv4.conf.all.accept_redirects = 0" is set in: "/etc/sysctl.d/60-
ipv4_sysctl.conf"
```

Note: If multiple lines with `net.ipv4.conf.all.accept_redirects =` are returned.

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Internal Only - General

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `net.ipv4.conf.all.accept_redirects` lines that are not `net.ipv4.conf.all.accept_redirects=0`.

Example script:

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

{
   l_option="net.ipv4.conf.all.accept_redirects" l_value="0"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

```
net.ipv4.conf.all.accept_redirects = 0
```

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv4.conf.all.accept_redirects = 0" \
>> /etc/sysctl.d/60-ipv4_sysctl.conf
```

Note: See the Configure IPv4 parameters subsection overview for important notes on UFW sysctl file interaction.

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv4.conf.all.accept_redirects = 1

Internal Only - General

References:

1. GRID: NIX-00000024
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-08-040279 | RULE ID: SV-244553r1017353 | CAT II
5. STIG ID: OL08-00-040279 | RULE ID: SV-248887r991589 | CAT II
6. STIG ID: RHEL-09-253015 | RULE ID: SV-257958r991589 | CAT II
7. STIG ID: ALMA-09-018610 | RULE ID: SV-269244r1050126 | CAT II
8. STIG ID: OL09-00-006020 | RULE ID: SV-271864r1092304 | CAT II
9. STIG ID: SLES-15-040330 | RULE ID: SV-235018r991589 | CAT II
10. CCI-000366

CIS Controls:

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

Internal Only - General

3.3.1.9 Ensure net.ipv4.conf.default.accept_redirects is configured (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.

`net.ipv4.conf.default.accept_redirects` controls accepting of all IPv4 ICMP redirected packets on a newly added network interface.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

Rationale:

ICMP redirect messages are packets that convey routing information and tell your host (acting as a router) to send packets via an alternate path. It is a way of allowing an outside routing device to update your system routing tables.

By setting `net.ipv4.conf.all.accept_redirects` to `0`, the system will not accept IPv4 ICMP redirect messages on a newly added network interface.

Audit:

Verify `net.ipv4.conf.default.accept_redirects` is set to `0`.

1. Run the following command to verify `net.ipv4.conf.default.accept_redirects` is set to `0` in the running configuration:

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

Verify output is:

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

2. Run the following script to verify `net.ipv4.conf.default.accept_redirects = 0` is set in a file being used by systemd sysctl to configure `net.ipv4.conf.default.accept_redirects`:

Internal Only - General

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

{
   l_parameter_name="net.ipv4.conf.default.accept_redirects"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv4.conf.default.accept_redirects = 0`.

Example output:

```
 - "net.ipv4.conf.default.accept_redirects = 0" is set in: "/etc/sysctl.d/60-
ipv4_sysctl.conf"
```

Note: If multiple lines with `net.ipv4.conf.default.accept_redirects =` are returned.

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Internal Only - General

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `net.ipv4.conf.default.accept_redirects` lines that are not `net.ipv4.conf.default.accept_redirects=0`.

Example script:

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

{
   l_option="net.ipv4.conf.default.accept_redirects" l_value="0"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

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

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv4.conf.default.accept_redirects = 0" \
>> /etc/sysctl.d/60-ipv4_sysctl.conf
```

Note: See the Configure IPv4 parameters subsection overview for important notes on UFW sysctl file interaction.

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv4.conf.default.accept_redirects = 1

Internal Only - General

References:

1. GRID: NIX-00000025
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-08-040209 | RULE ID: SV-244550r1017350 | CAT II
5. STIG ID: OL08-00-040209 | RULE ID: SV-248875r991589 | CAT II
6. STIG ID: RHEL-09-253040 | RULE ID: SV-257963r991589 | CAT II
7. STIG ID: ALMA-09-018610 | RULE ID: SV-269244r1050126 | CAT II
8. STIG ID: OL09-00-006025 | RULE ID: SV-271869r1092319 | CAT II
9. STIG ID: SLES-15-040340 | RULE ID: SV-235019r991589 | CAT II
10. CCI-000366

CIS Controls:

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

Internal Only - General

3.3.1.10 Ensure net.ipv4.conf.all.secure_redirects is configured (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.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

Rationale:

It is still possible for even known gateways to be compromised.

Setting `net.ipv4.conf.all.secure_redirects` to `0` the system will not accept any IPv4 routing table updates by possibly compromised known gateways.

Audit:

Verify `net.ipv4.conf.all.secure_redirects` is set to `0`.

1. Run the following command to verify `net.ipv4.conf.all.secure_redirects` is set to `0` in the running configuration:

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

Verify output is:

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

2. Run the following script to verify `net.ipv4.conf.all.secure_redirects = 0` is set in a file being used by systemd sysctl to configure `net.ipv4.conf.all.secure_redirects`:

Internal Only - General

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

{
   l_parameter_name="net.ipv4.conf.all.secure_redirects"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv4.conf.all.secure_redirects = 0`.

Example output:

```
 - "net.ipv4.conf.all.secure_redirects = 0" is set in: "/etc/sysctl.d/60-
ipv4_sysctl.conf"
```

Note: If multiple lines with `net.ipv4.conf.all.secure_redirects =` are returned.

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Internal Only - General

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `net.ipv4.conf.all.secure_redirects` lines that are not `net.ipv4.conf.all.secure_redirects=0`.

Example script:

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

{
   l_option="net.ipv4.conf.all.secure_redirects" l_value="0"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

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

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv4.conf.all.secure_redirects = 0" \
>> /etc/sysctl.d/60-ipv4_sysctl.conf
```

Note: See the Configure IPv4 parameters subsection overview for important notes on UFW sysctl file interaction.

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv4.conf.all.secure_redirects = 1

Internal Only - General

References:

1. GRID: NIX-00000026
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. CCI-000366

CIS Controls:

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

Internal Only - General

3.3.1.11 Ensure net.ipv4.conf.default.secure_redirects is configured (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.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

Rationale:

It is still possible for even known gateways to be compromised.

Setting `net.ipv4.conf.default.secure_redirects` to `0` the system the system will not accept IPv4 routing table updates by possibly compromised known gateways on a newly added network interface.

Audit:

Verify `net.ipv4.conf.default.secure_redirects` is set to `0`.

1. Run the following command to verify `net.ipv4.conf.default.secure_redirects` is set to `0` in the running configuration:

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

Verify output is:

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

2. Run the following script to verify `net.ipv4.conf.default.secure_redirects = 0` is set in a file being used by systemd sysctl to configure `net.ipv4.conf.default.secure_redirects`:

Internal Only - General

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

{
   l_parameter_name="net.ipv4.conf.default.secure_redirects"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv4.conf.default.secure_redirects = 0`.

Example output:

```
 - "net.ipv4.conf.default.secure_redirects = 0" is set in: "/etc/sysctl.d/60-
ipv4_sysctl.conf"
```

Note: If multiple lines with `net.ipv4.conf.default.secure_redirects =` are returned.

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Internal Only - General

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `net.ipv4.conf.default.secure_redirects` lines that are not `net.ipv4.conf.default.secure_redirects=0`.

Example script:

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

{
   l_option="net.ipv4.conf.default.secure_redirects" l_value="0"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

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

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv4.conf.default.secure_redirects = 0" \
>> /etc/sysctl.d/60-ipv4_sysctl.conf
```

Note: See the Configure IPv4 parameters subsection overview for important notes on UFW sysctl file interaction.

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv4.conf.default.secure_redirects = 1

Internal Only - General

References:

1. GRID: NIX-00000027
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. CCI-000366

CIS Controls:

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

Internal Only - General

3.3.1.12 Ensure net.ipv4.conf.all.rp_filter is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`net.ipv4.conf.all.rp_filter` controls reverse path filtering for IPv4 traffic on all network interfaces.

Reverse path filtering is a security mechanism in Linux that helps prevent IP spoofing and other network attacks. When enabled, it checks if the source IP address of an incoming packet is reachable through the same network interface it was received on. If the packet's source address is not reachable through the interface it arrived on, the packet is dropped.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

Rationale:

Setting `net.ipv4.conf.all.rp_filter` to `1` is a good way to deter attackers from sending your system bogus packets that cannot be responded to. One instance where this feature breaks down is if asymmetrical routing is employed. This would occur when using dynamic routing protocols (bgp, ospf, etc) on your system. If you are using asymmetrical routing on your system, you will not be able to enable this feature without breaking the routing.

Impact:

If you are using asymmetrical routing on your system, you will not be able to enable this feature without breaking the routing.

Audit:

Verify `net.ipv4.conf.all.rp_filter` is set to `1`.

1. Run the following command to verify `net.ipv4.conf.all.rp_filter` is set to `1` in the running configuration:

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

Verify output is:

Internal Only - General

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

2. Run the following script to verify `net.ipv4.conf.all.rp_filter = 1` is set in a file being used by systemd sysctl to configure `net.ipv4.conf.all.rp_filter`:

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

{
   l_parameter_name="net.ipv4.conf.all.rp_filter"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv4.conf.all.rp_filter = 1`.

Example output:

```
 - "net.ipv4.conf.all.rp_filter = 1" is set in: "/etc/sysctl.d/60-
ipv4_sysctl.conf"
```

Note: If multiple lines with `net.ipv4.conf.all.rp_filter =` are returned.

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Internal Only - General

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `net.ipv4.conf.all.rp_filter` lines that are not `net.ipv4.conf.all.rp_filter=1`.

Example script:

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

{
   l_option="net.ipv4.conf.all.rp_filter" l_value="1"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

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

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv4.conf.all.rp_filter = 1" \
>> /etc/sysctl.d/60-ipv4_sysctl.conf
```

Note: See the Configure IPv4 parameters subsection overview for important notes on UFW sysctl file interaction.

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv4.conf.all.rp_filter = 2

Internal Only - General

References:

1. GRID: NIX-00000028
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-08-040285 | RULE ID: SV-230549r1017311 | CAT II
5. STIG ID: OL08-00-040285 | RULE ID: SV-248893r991589 | CAT II
6. STIG ID: RHEL-09-253035 | RULE ID: SV-257962r991589 | CAT II
7. STIG ID: ALMA-09-019820 | RULE ID: SV-269255r1050137 | CAT II
8. STIG ID: OL09-00-006024 | RULE ID: SV-271868r1092316 | CAT II
9. CCI-000366

CIS Controls:

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

Internal Only - General

3.3.1.13 Ensure net.ipv4.conf.default.rp_filter is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`net.ipv4.conf.default.rp_filter` controls reverse path filtering for IPv4 traffic on newly added network interfaces.

Reverse path filtering is a security mechanism in Linux that helps prevent IP spoofing and other network attacks. When enabled, it checks if the source IP address of an incoming packet is reachable through the same network interface it was received on. If the packet's source address is not reachable through the interface it arrived on, the packet is dropped.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

Rationale:

Setting `net.ipv4.conf.default.rp_filter` to `1` is a good way to deter attackers from sending your system bogus packets that cannot be responded to. One instance where this feature breaks down is if asymmetrical routing is employed. This would occur when using dynamic routing protocols (bgp, ospf, etc) on your system. If you are using asymmetrical routing on your system, you will not be able to enable this feature without breaking the routing.

Impact:

If you are using asymmetrical routing on your system, you will not be able to enable this feature without breaking the routing.

Audit:

Verify `net.ipv4.conf.default.rp_filter` is set to `1`.

1. Run the following command to verify `net.ipv4.conf.default.rp_filter` is set to `1` in the running configuration:

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

Verify output is:

Internal Only - General

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

2. Run the following script to verify `net.ipv4.conf.default.rp_filter = 1` is set in a file being used by systemd sysctl to configure `net.ipv4.conf.default.rp_filter`:

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

{
   l_parameter_name="net.ipv4.conf.default.rp_filter"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv4.conf.default.rp_filter = 1`.

Example output:

```
 - "net.ipv4.conf.default.rp_filter = 1" is set in: "/etc/sysctl.d/60-
ipv4_sysctl.conf"
```

Note: If multiple lines with `net.ipv4.conf.default.rp_filter =` are returned.

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect

Internal Only - General

value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `net.ipv4.conf.default.rp_filter` lines that are not `net.ipv4.conf.default.rp_filter=1`.

Example script:

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

{
   l_option="net.ipv4.conf.default.rp_filter" l_value="1"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

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

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv4.conf.default.rp_filter = 1" \
>> /etc/sysctl.d/60-ipv4_sysctl.conf
```

Note: See the Configure IPv4 parameters subsection overview for important notes on UFW sysctl file interaction.

3. Run the following command to load all sysctl configuration filles:

Internal Only - General

```
# sysctl --system
```

Default Value:

net.ipv4.conf.default.rp_filter = 1

References:

1. GRID: NIX-00000029
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-09-253050 | RULE ID: SV-257965r991589 | CAT II
5. STIG ID: ALMA-09-019820 | RULE ID: SV-269255r1050137 | CAT II
6. STIG ID: OL09-00-006027 | RULE ID: SV-271871r1092325 | CAT II
7. CCI-000366

CIS Controls:

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

Internal Only - General

3.3.1.14 Ensure net.ipv4.conf.all.accept_source_route is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `accept_source_route` option causes network interfaces to accept packets with the Strict Source Route (SSR) or Loose Source Routing (LSR) option set.

`net.ipv4.conf.all.accept_source_route` controls accepting of all IPv4 packets with the SSR or LSR option set on all interfaces.

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.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

Rationale:

Setting `net.ipv4.conf.all.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:

Verify `net.ipv4.conf.all.accept_source_route` is set to `0`.

1. Run the following command to verify `net.ipv4.conf.all.accept_source_route` is set to `0` in the running configuration:

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

Verify output is:

Internal Only - General

```
net.ipv4.conf.all.accept_source_route = 0
```

2. Run the following script to verify `net.ipv4.conf.all.accept_source_route = 0` is set in a file being used by systemd sysctl to configure `net.ipv4.conf.all.accept_source_route`:

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

{
   l_parameter_name="net.ipv4.conf.all.accept_source_route"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv4.conf.all.accept_source_route = 0`.

Example output:

```
 - "net.ipv4.conf.all.accept_source_route = 0" is set in: "/etc/sysctl.d/60-
ipv4_sysctl.conf"
```

Note: If multiple lines with `net.ipv4.conf.all.accept_source_route =` are returned.

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect

Internal Only - General

value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `net.ipv4.conf.all.accept_source_route` lines that are not `net.ipv4.conf.all.accept_source_route=0`.

Example script:

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

{
   l_option="net.ipv4.conf.all.accept_source_route" l_value="0"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

```
net.ipv4.conf.all.accept_source_route = 0
```

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv4.conf.all.accept_source_route = 0" \
>> /etc/sysctl.d/60-ipv4_sysctl.conf
```

Note: See the Configure IPv4 parameters subsection overview for important notes on UFW sysctl file interaction.

3. Run the following command to load all sysctl configuration filles:

Internal Only - General

```
# sysctl --system
```

Default Value:

net.ipv4.conf.all.accept_source_route = 0

References:

1. GRID: NIX-00000030
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-08-040239 | RULE ID: SV-244551r1017351 | CAT II
5. STIG ID: OL08-00-040239 | RULE ID: SV-248879r991589 | CAT II
6. STIG ID: RHEL-09-253020 | RULE ID: SV-257959r991589 | CAT II
7. STIG ID: ALMA-09-020260 | RULE ID: SV-269259r1050141 | CAT II
8. STIG ID: OL09-00-006021 | RULE ID: SV-271865r1092307 | CAT II
9. STIG ID: SLES-15-040300 | RULE ID: SV-235014r991589 | CAT II
10. CCI-000366

CIS Controls:

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

Internal Only - General

3.3.1.15 Ensure net.ipv4.conf.default.accept_source_route is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `accept_source_route` option causes network interfaces to accept packets with the Strict Source Route (SSR) or Loose Source Routing (LSR) option set.

`net.ipv4.conf.default.accept_source_route` controls accepting of all IPv4 packets with the SSR or LSR option set on newly added network interfaces.

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.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

Rationale:

Setting `net.ipv4.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:

Verify `net.ipv4.conf.default.accept_source_route` is set to `0`.

1. Run the following command to verify `net.ipv4.conf.default.accept_source_route` is set to `0` in the running configuration:

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

Verify output is:

Internal Only - General

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

2. Run the following script to verify `net.ipv4.conf.default.accept_source_route = 0` is set in a file being used by systemd sysctl to configure `net.ipv4.conf.default.accept_source_route`:

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

{
   l_parameter_name="net.ipv4.conf.default.accept_source_route"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv4.conf.default.accept_source_route = 0`.

Example output:

```
 - "net.ipv4.conf.default.accept_source_route = 0" is set in:
"/etc/sysctl.d/60-ipv4_sysctl.conf"
```

Note: If multiple lines with `net.ipv4.conf.default.accept_source_route =` are returned.

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect

Internal Only - General

value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `net.ipv4.conf.default.accept_source_route` lines that are not `net.ipv4.conf.default.accept_source_route=0`.

Example script:

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

{
   l_option="net.ipv4.conf.default.accept_source_route" l_value="0"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

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

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv4.conf.default.accept_source_route = 0" \
>> /etc/sysctl.d/60-ipv4_sysctl.conf
```

Note: See the Configure IPv4 parameters subsection overview for important notes on UFW sysctl file interaction.

3. Run the following command to load all sysctl configuration filles:

Internal Only - General

```
# sysctl --system
```

Default Value:

net.ipv4.conf.default.accept_source_route = 1

References:

1. GRID: NIX-00000031
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-08-040249 | RULE ID: SV-244552r1017352 | CAT II
5. STIG ID: OL08-00-040249 | RULE ID: SV-248881r991589 | CAT II
6. STIG ID: RHEL-09-253045 | RULE ID: SV-257964r991589 | CAT II
7. STIG ID: ALMA-09-020260 | RULE ID: SV-269259r1050141 | CAT II
8. STIG ID: OL09-00-006026 | RULE ID: SV-271870r1092322 | CAT II
9. STIG ID: SLES-15-040320 | RULE ID: SV-235016r991589 | CAT II
10. CCI-000366

CIS Controls:

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

Internal Only - General

3.3.1.16 Ensure net.ipv4.conf.all.log_martians is configured (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.

`net.ipv4.conf.all.log_martians` controls if IPv4 packets with un-routable source addresses on any network interface is logged to the kernel log.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

Rationale:

Setting `net.ipv4.conf.all.log_martians` to `1` enables this feature. Logging these packets allows an administrator to investigate the possibility that an attacker is sending spoofed packets to their system.

Audit:

Verify `net.ipv4.conf.all.log_martians` is set to `1`.

1. Run the following command to verify `net.ipv4.conf.all.log_martians` is set to `1` in the running configuration:

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

Verify output is:

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

2. Run the following script to verify `net.ipv4.conf.all.log_martians = 1` is set in a file being used by systemd sysctl to configure `net.ipv4.conf.all.log_martians`:

Internal Only - General

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

{
   l_parameter_name="net.ipv4.conf.all.log_martians"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv4.conf.all.log_martians = 1`.

Example output:

```
 - "net.ipv4.conf.all.log_martians = 1" is set in: "/etc/sysctl.d/60-
ipv4_sysctl.conf"
```

Note: If multiple lines with `net.ipv4.conf.all.log_martians =` are returned.

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Internal Only - General

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `net.ipv4.conf.all.log_martians` lines that are not `net.ipv4.conf.all.log_martians=1`.

Example script:

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

{
   l_option="net.ipv4.conf.all.log_martians" l_value="1"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

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

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv4.conf.all.log_martians = 1" \
>> /etc/sysctl.d/60-ipv4_sysctl.conf
```

Note: See the Configure IPv4 parameters subsection overview for important notes on UFW sysctl file interaction.

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv4.conf.all.log_martians = 0

Internal Only - General

References:

1. GRID: NIX-00000032
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-09-253025 | RULE ID: SV-257960r991589 | CAT II
5. STIG ID: ALMA-09-019380 | RULE ID: SV-269251r1050133 | CAT II
6. STIG ID: OL09-00-006022 | RULE ID: SV-271866r1092310 | CAT II
7. CCI-000366

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.		●	●

Internal Only - General

3.3.1.17 Ensure net.ipv4.conf.default.log_martians is configured (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.

`net.ipv4.conf.default.log_martians` controls if IPv4 packets with un-routable source addresses on a newly added network interface is logged to the kernel log.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

Rationale:

Setting `net.ipv4.conf.default.log_martians` to `1` enables this feature. Logging these packets allows an administrator to investigate the possibility that an attacker is sending spoofed packets to their system.

Audit:

Verify `net.ipv4.conf.default.log_martians` is set to `1`.

1. Run the following command to verify `net.ipv4.conf.default.log_martians` is set to `1` in the running configuration:

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

Verify output is:

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

2. Run the following script to verify `net.ipv4.conf.default.log_martians = 1` is set in a file being used by systemd sysctl to configure `net.ipv4.conf.default.log_martians`:

Internal Only - General

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

{
   l_parameter_name="net.ipv4.conf.default.log_martians"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv4.conf.default.log_martians = 1`.

Example output:

```
 - "net.ipv4.conf.default.log_martians = 1" is set in: "/etc/sysctl.d/60-
ipv4_sysctl.conf"
```

Note: If multiple lines with `net.ipv4.conf.default.log_martians =` are returned.

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Internal Only - General

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `net.ipv4.conf.default.log_martians` lines that are not `net.ipv4.conf.default.log_martians=1`.

Example script:

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

{
   l_option="net.ipv4.conf.default.log_martians" l_value="1"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

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

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv4.conf.default.log_martians = 1" \
>> /etc/sysctl.d/60-ipv4_sysctl.conf
```

Note: See the Configure IPv4 parameters subsection overview for important notes on UFW sysctl file interaction.

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv4.conf.default.log_martians = 0

Internal Only - General

References:

1. GRID: NIX-00000033
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-09-253030 | RULE ID: SV-257961r991589 | CAT II
5. STIG ID: ALMA-09-019380 | RULE ID: SV-269251r1050133 | CAT II
6. STIG ID: OL09-00-006023 | RULE ID: SV-271867r1092313 | CAT II
7. CCI-000366

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.		●	●

Internal Only - General

3.3.1.18 Ensure net.ipv4.tcp_syncookies is configured (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.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

Rationale:

Attackers use SYN flood attacks to perform a denial of service attacked on a system by sending many SYN packets without completing the three way handshake. This will quickly use up slots in the kernel's half-open connection queue and prevent legitimate connections from succeeding. Setting `net.ipv4.tcp_syncookies` to `1` enables SYN cookies, allowing the system to keep accepting valid connections, even if under a denial of service attack.

Audit:

Verify `net.ipv4.tcp_syncookies` is set to `1`.

1. Run the following command to verify `net.ipv4.tcp_syncookies` is set to `1` in the running configuration:

```
# sysctl net.ipv4.tcp_syncookies
```

Verify output is:

```
net.ipv4.tcp_syncookies = 1
```

2. Run the following script to verify `net.ipv4.tcp_syncookies = 1` is set in a file being used by systemd sysctl to configure `net.ipv4.tcp_syncookies`:

Internal Only - General

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

{
   l_parameter_name="net.ipv4.tcp_syncookies"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv4.tcp_syncookies = 1`.

Example output:

```
 - "net.ipv4.tcp_syncookies = 1" is set in: "/etc/sysctl.d/60-
ipv4_sysctl.conf"
```

Note: If multiple lines with `net.ipv4.tcp_syncookies =` are returned.

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `net.ipv4.tcp_syncookies` lines that are not `net.ipv4.tcp_syncookies=1`.

Internal Only - General

Example script:

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

{
   l_option="net.ipv4.tcp_syncookies" l_value="1"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

```
net.ipv4.tcp_syncookies = 1
```

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv4.tcp_syncookies = 1" >> /etc/sysctl.d/60-
ipv4_sysctl.conf
```

Note: See the Configure IPv4 parameters subsection overview for important notes on UFW sysctl file interaction.

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv4.tcp_syncookies = 1

References:

1. GRID: NIX-00000034
2. NIST SP 800-53 :: SC-5 (2) | NIST SP 800-53 :: SC-5 a
3. NIST SP 800-53A :: SC-5 (2).1
4. STIG ID: RHEL-09-253010 | RULE ID: SV-257957r1045009 | CAT II

Internal Only - General

5. STIG ID: ALMA-09-042480 | RULE ID: SV-269435r1050318 | CAT II
6. STIG ID: OL09-00-006050 | RULE ID: SV-271884r1092364 | CAT II
7. STIG ID: UBTU-22-253010 | RULE ID: SV-260522r1069097 | CAT II
8. STIG ID: UBTU-24-600190 | RULE ID: SV-270753r1066748 | CAT II
9. STIG ID: SLES-15-010310 | RULE ID: SV-234829r958528 | CAT II
10. CCI-001095 | CCI-002385

CIS Controls:

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

Internal Only - General

3.3.1.19 Ensure net.ipv4.conf.all.route_localnet is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `net.ipv4.conf.all.route_localnet` flag enables the use of `127/8` for local routing purposes and tells the system do not consider loopback addresses as martian source or destination while routing.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

Rationale:

Setting `net.ipv4.conf.all.route_localnet` to `0` enforces network isolation by preventing external interfaces from routing traffic to loopback addresses.

Audit:

Verify `net.ipv4.conf.all.route_localnet` is set to `0`.

1. Run the following command to verify that `net.ipv4.conf.all.route_localnet` is set to `0` in the running configuration:

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

Verify output is:

```
net.ipv4.conf.all.route_localnet = 0
```

2. Run the following script to verify `net.ipv4.conf.all.route_localnet = 0` is set in a file being used by systemd sysctl to configure `net.ipv4.conf.all.route_localnet`:

Internal Only - General

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

{
   l_parameter_name="net.ipv4.conf.all.route_localnet"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes: `net.ipv4.conf.all.route_localnet = 0`

Example output

```
- "net.ipv4.conf.all.route_localnet = 0" is set in: "/etc/sysctl.d/60-
ipv4_sysctl.conf"
```

Remediation:

Review all files being used by systemd sysctl and comment out or remove all `net.ipv4.conf.all.route_localnet` lines that are not `net.ipv4.conf.all.route_localnet = 0`.

Example script

Internal Only - General

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

{
   l_option="net.ipv4.conf.all.route_localnet" l_value="0"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf` and edit or add the following line: `net.ipv4.conf.all.route_localnet = 0`

Example

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv4.conf.all.route_localnet = 0" \
>> /etc/sysctl.d/60-ipv4_sysctl.conf
```

Note: See the Configure IPv4 parameters subsection overview for important notes on UFW sysctl file interaction.

Run the following command to load all sysctl configuration files:

```
# sysctl --system
```

References:

1. GRID: NIX-00000386

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 Uninstall or Disable Unnecessary Services on Enterprise Assets and Software Uninstall or disable unnecessary services on enterprise assets and software, such as an unused file sharing service, web application module, or service function.		●	●

Internal Only - General

Controls Version	Control	IG 1	IG 2	IG 3
v7	9.2 Ensure Only Approved Ports, Protocols and Services Are Running Ensure that only network ports, protocols, and services listening on a system with validated business needs, are running on each system.		●	●

Internal Only - General

3.3.2 Configure IPv6 parameters

net.ipv6 kernel parameters control various aspects of the IPv6 network stack. These parameters effect functions to include routing, TCP, and security.

Note: If IPv6 has been disabled on the system, the Recommendations in this subsection are not applicable and should be skipped.

The following script creates and runs a function called f_ipv6_chk that may be used to determine if IPv6 has been disabled on the system:

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

f_ipv6_chk()
{
   a_output=()
   ! grep -Pqs -- '^\h*0\b' /sys/module/ipv6/parameters/disable && \
   a_output+=(" - IPv6 is disabled")
   if sysctl net.ipv6.conf.all.disable_ipv6 | \
   grep -Pqs -- "^\h*net\.ipv6\.conf\.all\.disable_ipv6\h*=\h*1\b" && \
   sysctl net.ipv6.conf.default.disable_ipv6 | \
   grep -Pqs -- "^\h*net\.ipv6\.conf\.default\.disable_ipv6\h*=\h*1\b"; then
      [ "${#a_output[@]}" -le 0 ] && a_output+=(" - IPv6 is disabled")
   fi
   [ "${#a_output[@]}" -le 0 ] && a_output+=(" - IPv6 is not disabled")
   printf '%s\n' "" "${a_output[@]}" ""
}
f_ipv6_chk
```

Router Applicability

The following recommendations in this subsection are not applicable to systems configured as network routers, as these functions are required for correct router operation:

- 3.3.2.1 — Ensure net.ipv6.conf.all.forwarding is configured
- 3.3.2.2 — Ensure net.ipv6.conf.default.forwarding is configured
- 3.3.2.3 — Ensure net.ipv6.conf.all.accept_redirects is configured
- 3.3.2.4 — Ensure net.ipv6.conf.default.accept_redirects is configured

Routers require IPv6 packet forwarding to be enabled to route traffic between interfaces, and must be able to receive ICMP redirect messages to update routing tables. Recommendations 3.3.2.5 through 3.3.2.8 apply to all systems, including network routers.

Parameter Interdependencies

net.ipv6.conf.all.* parameters take precedence over net.ipv6.conf.default.* parameters for all currently active interfaces. The conf.default.* parameters establish the baseline applied to interfaces added after boot. Both the all and default variants of each parameter must be configured to ensure consistent behavior across all current and future network interfaces.

Internal Only - General

UFW Sysctl File Interaction

On systems with Uncomplicated Firewall (UFW) installed, `/etc/ufw/sysctl.conf` may appear as the first file identified during the audit procedure. UFW maintains its own sysctl configuration file (referenced by the `IPT_SYSCTL` setting in `/etc/default/ufw`) and this file is processed earlier in the sysctl load order than files in `/etc/sysctl.d/`. As a result:

- If an incorrect value for any parameter covered by this subsection is present in `/etc/ufw/sysctl.conf`, the first remediation step will comment it out.
- Future UFW package updates may rewrite or re-enable entries in `/etc/ufw/sysctl.conf`, potentially overriding values set during remediation. If this occurs, the UFW sysctl file will take precedence over values written to `/etc/sysctl.d/` because it is processed earlier in the load order.

On systems where UFW is installed, periodically verify that the settings in this subsection persist after UFW package updates.

Internal Only - General

3.3.2.1 Ensure net.ipv6.conf.all.forwarding is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `net.ipv6.conf.all.forwarding` flag tells the system whether it can forward IPv6 packets or not.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

> **Note: If IPv6 has been disabled, or this system is a router, this recommendation is not applicable.**

Rationale:

Routing protocol daemons are typically used on routers to exchange network topology information with other routers. If this software is used when not required, system network information may be unnecessarily transmitted across the network.

Setting `net.ipv6.conf.all.forwarding` to `0` ensures that a system with multiple interfaces (for example, a hard proxy), will not be able to forward IPv6 packets.

Impact:

IP forwarding is required on systems configured to act as a router. If these parameters are disabled, the system will not be able to perform as a router.

Many Cloud Service Provider (CSP) hosted systems require IP forwarding to be enabled. If the system is running on a CSP platform, this requirement should be reviewed before disabling IP forwarding.

Audit:

> **- IF IPv6 is enabled on the system -**

Verify `net.ipv6.conf.all.forwarding` is set to `0`.

1. Run the following command to verify `net.ipv6.conf.all.forwarding` is set to `0` in the running configuration:

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

Verify output is:

Internal Only - General

```
net.ipv6.conf.all.forwarding = 0
```

2. Run the following script to verify `net.ipv6.conf.all.forwarding = 0` is set in a file being used by systemd sysctl to configure `net.ipv6.conf.all.forwarding`:

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

{
   l_parameter_name="net.ipv6.conf.all.forwarding"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv6.conf.all.forwarding = 0`.

Example output:

```
 - "net.ipv6.conf.all.forwarding = 0" is set in: "/etc/sysctl.d/60-
ipv6_sysctl.conf"
```

Note: If multiple lines with `net.ipv6.conf.all.forwarding =` are returned.

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect

Internal Only - General

value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `net.ipv6.conf.all.forwarding` lines that are not `net.ipv6.conf.all.forwarding=0`.

Example script:

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

{
   l_option="net.ipv6.conf.all.forwarding" l_value="0"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

```
net.ipv6.conf.all.forwarding = 0
```

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv6.conf.all.forwarding = 0" >> /etc/sysctl.d/60-
ipv6_sysctl.conf
```

Note: If the UFW file was the first file listed in the audit, the entry will be commented out as part of the first step, however updating Uncomplicated Firewall (UFW) may update this change. In this case the updated entry will supersede the entry being created as part of this step.

3. Run the following command to load all sysctl configuration filles:

Internal Only - General

```
# sysctl --system
```

Default Value:

net.ipv6.conf.all.forwarding = 0

References:

1. GRID: NIX-00000035
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-08-040260 | RULE ID: SV-230540r1017302 | CAT II
5. STIG ID: OL08-00-040260 | RULE ID: SV-248883r991589 | CAT II
6. STIG ID: RHEL-09-254025 | RULE ID: SV-257974r991589 | CAT II
7. STIG ID: ALMA-09-019160 | RULE ID: SV-269249r1050131 | CAT II
8. STIG ID: OL09-00-006043 | RULE ID: SV-271880r1092352 | CAT II
9. STIG ID: SLES-15-040381 | RULE ID: SV-235025r991589 | CAT II
10. CI-000366

CIS Controls:

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

Internal Only - General

3.3.2.2 Ensure net.ipv6.conf.default.forwarding is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `net.ipv6.conf.default.forwarding` flag tells the system whether it can forward IPv6 on the default or newly added network interface.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

> **Note: If IPv6 has been disabled, or this system is a router, this recommendation is not applicable.**

Rationale:

Routing protocol daemons are typically used on routers to exchange network topology information with other routers. If this software is used when not required, system network information may be unnecessarily transmitted across the network.

Impact:

IP forwarding is required on systems configured to act as a router. If these parameters are disabled, the system will not be able to perform as a router.

Many Cloud Service Provider (CSP) hosted systems require IP forwarding to be enabled. If the system is running on a CSP platform, this requirement should be reviewed before disabling IP forwarding.

Audit:

> **- IF IPv6 is enabled on the system -**

Verify `net.ipv6.conf.default.forwarding` is set to `0`.

1. Run the following command to verify `net.ipv6.conf.default.forwarding` is set to `0` in the running configuration:

```
# sysctl net.ipv6.conf.default.forwarding
```

Verify output is:

Internal Only - General

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

2. Run the following script to verify `net.ipv6.conf.default.forwarding = 0` is set in a file being used by systemd sysctl to configure `net.ipv6.conf.default.forwarding`:

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

{
   l_parameter_name="net.ipv6.conf.default.forwarding"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv6.conf.default.forwarding = 0`.

Example output:

```
 - "net.ipv6.conf.default.forwarding = 0" is set in: "/etc/sysctl.d/60-
ipv6_sysctl.conf"
```

Note: If multiple lines with `net.ipv6.conf.default.forwarding =` are returned.

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect

Internal Only - General

value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `net.ipv6.conf.default.forwarding` lines that are not `net.ipv6.conf.default.forwarding=0`.

Example script:

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

{
   l_option="net.ipv6.conf.default.forwarding" l_value="0"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

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

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv6.conf.default.forwarding = 0" >>
/etc/sysctl.d/60-ipv6_sysctl.conf
```

Note: If the UFW file was the first file listed in the audit, the entry will be commented out as part of the first step, however updating Uncomplicated Firewall (UFW) may update this change. In this case the updated entry will supersede the entry being created as part of this step.

3. Run the following command to load all sysctl configuration filles:

Internal Only - General

```
# sysctl --system
```

Default Value:

net.ipv6.conf.default.forwarding = 0

References:

1. GRID: NIX-00000036
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: ALMA-09-019160 | RULE ID: SV-269249r1050131 | CAT II
5. STIG ID: SLES-15-040382 | RULE ID: SV-235026r991589 | CAT II
6. CI-000366

CIS Controls:

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

Internal Only - General

3.3.2.3 Ensure net.ipv6.conf.all.accept_redirects is configured (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.

`net.ipv6.conf.all.accept_redirects` controls accepting of all IPv6 ICMP redirected packets on all interfaces.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

> **Note: If IPv6 has been disabled, or this system is a router, this recommendation is not applicable.**

Rationale:

ICMP redirect messages are packets that convey routing information and tell your host (acting as a router) to send packets via an alternate path. It is a way of allowing an outside routing device to update your system routing tables.

By setting `net.ipv6.conf.all.accept_redirects` to `0`, the system will not accept any IPv6 ICMP redirect messages, and therefore, won't allow outsiders to update the system's routing tables.

Audit:

> **- IF IPv6 is enabled on the system -**

Verify `net.ipv6.conf.all.accept_redirects` is set to `0`.

1. Run the following command to verify `net.ipv6.conf.all.accept_redirects` is set to `0` in the running configuration:

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

Verify output is:

Internal Only - General

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

2. Run the following script to verify `net.ipv6.conf.all.accept_redirects = 0` is set in a file being used by systemd sysctl to configure `net.ipv6.conf.all.accept_redirects`:

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

{
   l_parameter_name="net.ipv6.conf.all.accept_redirects"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv6.conf.all.accept_redirects = 0`.

Example output:

```
 - "net.ipv6.conf.all.accept_redirects = 0" is set in: "/etc/sysctl.d/60-
ipv6_sysctl.conf"
```

Note: If multiple lines with `net.ipv6.conf.all.accept_redirects =` are returned.

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect

Internal Only - General

value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `net.ipv6.conf.all.accept_redirects` lines that are not `net.ipv6.conf.all.accept_redirects=0`.

Example script:

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

{
   l_option="net.ipv6.conf.all.accept_redirects" l_value="0"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

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

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv6.conf.all.accept_redirects = 0" \
>> /etc/sysctl.d/60-ipv6_sysctl.conf
```

Note: If the UFW file was the first file listed in the audit, the entry will be commented out as part of the first step, however updating Uncomplicated Firewall (UFW) may update this change. In this case the updated entry will supersede the entry being created as part of this step.

3. Run the following command to load all sysctl configuration filles:

Internal Only - General

```
# sysctl --system
```

Default Value:

net.ipv6.conf.all.accept_redirects = 1

References:

1. GRID: NIX-00000037
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-08-040280 | RULE ID: SV-230544r1017306 | CAT II
5. STIG ID: OL08-00-040280 | RULE ID: SV-248888r991589 | CAT II
6. STIG ID: RHEL-09-254015 | RULE ID: SV-257972r991589 | CAT II
7. STIG ID: ALMA-09-018610 | RULE ID: SV-269244r1050126 | CAT II
8. STIG ID: OL09-00-006041 | RULE ID: SV-271878r1092346 | CAT II
9. STIG ID: SLES-15-040341 | RULE ID: SV-235020r991589 | CAT II
10. CCI-000366

CIS Controls:

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

Internal Only - General

3.3.2.4 Ensure net.ipv6.conf.default.accept_redirects is configured (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.

`net.ipv6.conf.default.accept_redirects` controls accepting of all IPv6 ICMP redirected packets on a newly added network interface.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

> **Note: If IPv6 has been disabled, or this system is a router, this recommendation is not applicable.**

Rationale:

ICMP redirect messages are packets that convey routing information and tell your host (acting as a router) to send packets via an alternate path. It is a way of allowing an outside routing device to update your system routing tables.

By setting `net.ipv6.conf.default.accept_redirects` to `0`, the system will not accept IPv6 ICMP redirect messages on a newly added network interface.

Audit:

> **- IF IPv6 is enabled on the system -**

Verify `net.ipv6.conf.default.accept_redirects` is set to `0`.

1. Run the following command to verify `net.ipv6.conf.default.accept_redirects` is set to `0` in the running configuration:

```
# sysctl net.ipv6.conf.default.accept_redirects
```

Verify output is:

Internal Only - General

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

2. Run the following script to verify `net.ipv6.conf.default.accept_redirects = 0` is set in a file being used by systemd sysctl to configure `net.ipv6.conf.default.accept_redirects`:

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

{
   l_parameter_name="net.ipv6.conf.default.accept_redirects"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv6.conf.default.accept_redirects = 0`.

Example output:

```
 - "net.ipv6.conf.default.accept_redirects = 0" is set in: "/etc/sysctl.d/60-
ipv6_sysctl.conf"
```

Note: If multiple lines with `net.ipv6.conf.default.accept_redirects =` are returned.

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect

Internal Only - General

value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `net.ipv6.conf.default.accept_redirects` lines that are not `net.ipv6.conf.default.accept_redirects=0`.

Example script:

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

{
   l_option="net.ipv6.conf.default.accept_redirects" l_value="0"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

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

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv6.conf.default.accept_redirects = 0" \
>> /etc/sysctl.d/60-ipv6_sysctl.conf
```

Note: If the UFW file was the first file listed in the audit, the entry will be commented out as part of the first step, however updating Uncomplicated Firewall (UFW) may update this change. In this case the updated entry will supersede the entry being created as part of this step.

3. Run the following command to load all sysctl configuration filles:

Internal Only - General

```
# sysctl --system
```

Default Value:

net.ipv6.conf.default.accept_redirects = 1

References:

1. GRID: NIX-00000038
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-08-040210 | RULE ID: SV-230535r1017297 | CAT II
5. STIG ID: OL08-00-040210 | RULE ID: SV-248876r991589 | CAT II
6. STIG ID: RHEL-09-254035 | RULE ID: SV-257976r991589 | CAT II
7. STIG ID: ALMA-09-018610 | RULE ID: SV-269244r1050126 | CAT II
8. STIG ID: OL09-00-006045 | RULE ID: SV-271882r1092358 | CAT II
9. STIG ID: SLES-15-040350 | RULE ID: SV-235021r991589 | CAT II
10. CCI-000366

CIS Controls:

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

Internal Only - General

3.3.2.5 Ensure net.ipv6.conf.all.accept_source_route is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `accept_source_route` option causes network interfaces to accept packets with the Strict Source Route (SSR) or Loose Source Routing (LSR) option set.

`net.ipv6.conf.all.accept_source_route` controls accepting of all IPv6 packets with the SSR or LSR option set on all interfaces.

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.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

> **Note: If IPv6 has been disabled this recommendation is not applicable.**

Rationale:

Setting `net.ipv6.conf.all.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:

> **- IF IPv6 is enabled on the system -**

Verify `net.ipv6.conf.all.accept_source_route` is set to `0`.

1. Run the following command to verify `net.ipv6.conf.all.accept_source_route` is set to `0` in the running configuration:

Internal Only - General

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

Verify output is:

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

2. Run the following script to verify `net.ipv6.conf.all.accept_source_route = 0` is set in a file being used by systemd sysctl to configure `net.ipv6.conf.all.accept_source_route`:

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

{
   l_parameter_name="net.ipv6.conf.all.accept_source_route"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv4.conf.default.forwarding = 0`.

Example output:

```
 - "net.ipv6.conf.all.accept_source_route = 0" is set in: "/etc/sysctl.d/60-
ipv6_sysctl.conf"
```

Note: If multiple lines with `net.ipv6.conf.all.accept_source_route =` are returned.

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.

Internal Only - General

- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `net.ipv6.conf.all.accept_source_route` lines that are not `net.ipv6.conf.all.accept_source_route=0`.

Example script:

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

{
   l_option="net.ipv6.conf.all.accept_source_route" l_value="0"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

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

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv6.conf.all.accept_source_route = 0" \
>> /etc/sysctl.d/60-ipv6_sysctl.conf
```

Note: If the UFW file was the first file listed in the audit, the entry will be commented out as part of the first step, however updating Uncomplicated Firewall (UFW) may update this change. In this case the updated entry will supersede the entry being created as part of this step.

Internal Only - General

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv6.conf.all.accept_source_route = 0

References:

1. GRID: NIX-00000039
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-08-040240 | RULE ID: SV-230538r1017300 | CAT II
5. STIG ID: OL08-00-040240 | RULE ID: SV-248880r991589 | CAT II
6. STIG ID: RHEL-09-254020 | RULE ID: SV-257973r991589 | CAT II
7. STIG ID: ALMA-09-020260 | RULE ID: SV-269259r1050141 | CAT II
8. STIG ID: OL09-00-006042 | RULE ID: SV-271879r1092349 | CAT II
9. STIG ID: SLES-15-040310 | RULE ID: SV-235015r991589 | CAT II
10. CCI-000366

CIS Controls:

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

Internal Only - General

3.3.2.6 Ensure net.ipv6.conf.default.accept_source_route is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `accept_source_route` option causes network interfaces to accept packets with the Strict Source Route (SSR) or Loose Source Routing (LSR) option set.

`net.ipv6.conf.default.accept_source_route` controls accepting of all IPv6 packets with the SSR or LSR option set on newly added network interfaces.

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.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

> **Note: If IPv6 has been disabled this recommendation is not applicable.**

Rationale:

Setting `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:

> **- IF IPv6 is enabled on the system -**

Verify `net.ipv6.conf.default.accept_source_route` is set to `0`.

1. Run the following command to verify `net.ipv6.conf.default.accept_source_route` is set to `0` in the running configuration:

Internal Only - General

```
# sysctl net.ipv6.conf.default.accept_source_route
```

Verify output is:

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

2. Run the following script to verify `net.ipv6.conf.default.accept_source_route = 0` is set in a file being used by systemd sysctl to configure `net.ipv6.conf.default.accept_source_route`:

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

{
   l_parameter_name="net.ipv6.conf.default.accept_source_route"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv6.conf.default.accept_source_route = 0`.

Example output:

```
 - "net.ipv6.conf.default.accept_source_route = 0" is set in:
"/etc/sysctl.d/60-ipv6_sysctl.conf"
```

Note: If multiple lines with `net.ipv6.conf.default.accept_source_route =` are returned.

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a

Internal Only - General

`.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.

- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `net.ipv6.conf.default.accept_source_route` lines that are not `net.ipv6.conf.default.accept_source_route=0`.

Example script:

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

{
   l_option="net.ipv6.conf.default.accept_source_route" l_value="0"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

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

Example:

Internal Only - General

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv6.conf.default.accept_source_route = 0" \
>> /etc/sysctl.d/60-ipv6_sysctl.conf
```

Note: If the UFW file was the first file listed in the audit, the entry will be commented out as part of the first step, however updating Uncomplicated Firewall (UFW) may update this change. In this case the updated entry will supersede the entry being created as part of this step.

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv6.conf.default.accept_source_route = 0

References:

1. GRID: NIX-00000040
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-08-040250 | RULE ID: SV-230539r1017301 | CAT II
5. STIG ID: OL08-00-040250 | RULE ID: SV-248882r991589 | CAT II
6. STIG ID: RHEL-09-254040 | RULE ID: SV-257977r991589 | CAT II
7. STIG ID: ALMA-09-020260 | RULE ID: SV-269259r1050141 | CAT II
8. STIG ID: OL09-00-006046 | RULE ID: SV-271883r1092361 | CAT II
9. STIG ID: SLES-15-040321 | RULE ID: SV-235017r991589 | CAT II
10. CCI-000366

CIS Controls:

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

Internal Only - General

3.3.2.7 Ensure net.ipv6.conf.all.accept_ra is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Routers periodically multicast Router Advertisement messages to announce their availability and convey information to neighboring nodes that enable them to be automatically configured on the network.

`net.ipv6.conf.all.accept_ra` controls accepting any IPv6 RA messages on any network device.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

> **Note: If IPv6 has been disabled, this recommendation is not applicable.**

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:

> **- IF IPv6 is enabled on the system -**

Verify `net.ipv6.conf.all.accept_ra` is set to `0`.

1. Run the following command to verify `net.ipv6.conf.all.accept_ra` is set to `0` in the running configuration:

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

Verify output is:

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

2. Run the following script to verify `net.ipv6.conf.all.accept_ra = 0` is set in a file being used by systemd sysctl to configure `net.ipv6.conf.all.accept_ra`:

Internal Only - General

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

{
   l_parameter_name="net.ipv6.conf.all.accept_ra"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv6.conf.all.accept_ra = 0`.

Example output:

```
 - "net.ipv6.conf.all.accept_ra = 0" is set in: "/etc/sysctl.d/60-
ipv6_sysctl.conf"
```

Note: If multiple lines with `net.ipv6.conf.all.accept_ra =` are returned.

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Internal Only - General

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `net.ipv6.conf.all.accept_ra` lines that are not `net.ipv6.conf.all.accept_ra=0`.

Example script:

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

{
   l_option="net.ipv6.conf.all.accept_ra" l_value="0"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

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

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv6.conf.all.accept_ra = 0" >> /etc/sysctl.d/60-
ipv6_sysctl.conf
```

Note: If the UFW file was the first file listed in the audit, the entry will be commented out as part of the first step, however updating Uncomplicated Firewall (UFW) may update this change. In this case the updated entry will supersede the entry being created as part of this step.

3. Run the following command to load all sysctl configuration filles:

Internal Only - General

```
# sysctl --system
```

Default Value:

net.ipv6.conf.all.accept_ra = 1

References:

1. GRID: NIX-00000041
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-08-040261 | RULE ID: SV-230541r1017303 | CAT II
5. STIG ID: OL08-00-040261 | RULE ID: SV-248884r991589 | CAT II
6. STIG ID: RHEL-09-254010 | RULE ID: SV-257971r991589 | CAT II
7. STIG ID: ALMA-09-018500 | RULE ID: SV-269243r1050125 | CAT II
8. STIG ID: OL09-00-006040 | RULE ID: SV-271877r1092343 | CAT II
9. CCI-000366

CIS Controls:

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

Internal Only - General

3.3.2.8 Ensure net.ipv6.conf.default.accept_ra is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Routers periodically Multicast Router Advertisement (RA) messages to announce their availability and convey information to neighboring nodes that enable them to be automatically configured on the network.

`net.ipv6.conf.default.accept_ra` controls accepting IPv6 RA messages on newly added network devices.

More information about the kernel parameter configuration files, their location, and load preference is available in the "Configure Network Kernel Parameters" section overview.

> **Note: If IPv6 has been disabled, this recommendation is not applicable.**

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:

> **- IF IPv6 is enabled on the system -**

Verify `net.ipv6.conf.default.accept_ra` is set to `0`.

1. Run the following command to verify `net.ipv6.conf.default.accept_ra` is set to `0` in the running configuration:

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

Verify output is:

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

2. Run the following script to verify `net.ipv6.conf.default.accept_ra = 0` is set in a file being used by systemd sysctl to configure `net.ipv6.conf.default.accept_ra`:

Internal Only - General

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

{
   l_parameter_name="net.ipv6.conf.default.accept_ra"
   l_grep="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      l_opt="$(grep -Psio '^\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" | tail -n
1)"
      l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_parameter_name = $l_option_value\" is set in:
\"$l_file\"")
   done
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Verify output includes `net.ipv6.conf.default.accept_ra = 0`.

Example output:

```
 - "net.ipv6.conf.default.accept_ra = 0" is set in: "/etc/sysctl.d/60-
ipv6_sysctl.conf"
```

Note: If multiple lines with `net.ipv6.conf.default.accept_ra =` are returned.

- The first line includes the value being used by `systemd sysctl`. If this is a correct value, this is considered a passing state. If the file listed is not in the `/etc/sysctl.d/` directory, it is highly recommended to follow the remediation procedure to create a `.conf` file in the `/etc/sysctl.d/` directory with the correct setting to prevent a potential change due to an update to the system.
- Any files in the `/etc/sysctl.d/` directory that include an incorrect value should be modified to comment out or change the incorrect value to minimize the potential of the incorrect value being used by `systemd sysctl` due to system configuration changes.

Internal Only - General

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `net.ipv6.conf.default.accept_ra` lines that are not `net.ipv6.conf.default.accept_ra=0`.

Example script:

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

{
   l_option="net.ipv6.conf.default.accept_ra" l_value="0"
   l_grep="${l_option//./(\\.|\\/)}" a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   l_ufw_file="$([ -f /etc/default/ufw ] && \
   awk -F= '/^\s*IPT_SYSCTL=/ {print $2}' /etc/default/ufw)"
   [ -f "$(readlink -e "$l_ufw_file")" ] && \
   a_files+=("$l_ufw_file"); a_files+=("/etc/sysctl.conf")
   while IFS= read -r l_fname; do
      l_file="$(readlink -e "${l_fname//# /}")"
      [ -n "$l_file" ] && ! grep -Psiq -- '(^|\h+)'"$l_file"'\b' \
      <<< "${a_files[*]}" && a_files+=("$l_file")
   done < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Psio -- '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Psio -- '\h*'"$l_grep"'\h*=\h*\H+\b' "$l_file" \
      | grep -Psivq -- '^\h*'"$l_grep"'\h*=\h*'"$l_value"'\b' && \
      sed -ri '/^\s*'"$l_grep"'\s*=\s*(1[0-9]*)/s/^/# /' "$l_file"
   done
}
```

2. Create or edit a file in the `/etc/sysctl.d/` directory ending in `.conf`and edit or add the following line:

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

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv6.conf.default.accept_ra = 0" >> /etc/sysctl.d/60-
ipv6_sysctl.conf
```

Note: If the UFW file was the first file listed in the audit, the entry will be commented out as part of the first step, however updating Uncomplicated Firewall (UFW) may update this change. In this case the updated entry will supersede the entry being created as part of this step.

3. Run the following command to load all sysctl configuration filles:

Internal Only - General

```
# sysctl --system
```

Default Value:

net.ipv6.conf.default.accept_ra = 1

References:

1. GRID: NIX-00000042
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-08-040262 | RULE ID: SV-230542r1017304 | CAT II
5. STIG ID: OL08-00-040262 | RULE ID: SV-248885r991589 | CAT II
6. STIG ID: RHEL-09-254030 | RULE ID: SV-257975r991589 | CAT II
7. STIG ID: ALMA-09-018500 | RULE ID: SV-269243r1050125 | CAT II
8. STIG ID: OL09-00-006044 | RULE ID: SV-271881r1092355 | CAT II
9. CCI-000366

CIS Controls:

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

4 Host Based Firewall

A firewall is a set of rules. When a data packet moves into or out of a protected network space, its contents (in particular, information about its origin, target, and the protocol it plans to use) are tested against the firewall rules to see if it should be allowed through.

To provide a Host Based Firewall, the Linux kernel includes support for NFTables.

NFTables is subsystem of the Linux kernel providing filtering and classification of network packets/datagrams/frames. NFTables replaces parts of Netfilter, while keeping and reusing most of it. NFTables utilizes the building blocks of the Netfilter infrastructure, such as the existing hooks into the networking stack, connection tracking system, userspace queueing component, and logging subsystem.

Internal Only - General

In order to configure firewall rules for NFTables, a firewall utility needs to be installed. Guidance has been included for Uncomplicated Firewall (UFW) which provides firewall features by acting as a front-end for the Linux kernel's netfilter framework via the NFTables backend, acting as an alternative to the `nft` command line program. UFW supports both IPv4 and IPv6 networks.

Important: This Benchmark includes guidance for configuring `NFTables` though the `UFW` front end utility. **If your organization is using a different Host Based Firewall, or a different utility to configure the Host Based Firewall**, we recommend verifying that your site's configuration meets the Recommendations in this section, otherwise **this section may be skipped**.

Note:

- Only **one** method should be used to configure a firewall on the system. Use of more than one method could produce unexpected results.
- NFTables is the default packet filtering framework. Though IPTables is still supported, this support is done indirectly through UFW's ability to translate IPTables commands to NFTables.
- This section is intended only to ensure the resulting firewall rules are in place.

Internal Only - General

4.1 Configure Uncomplicated Firewall

In order to configure durable firewall rules a firewall utility needs to be installed. Uncomplicated Firewall (UFW) is the recommended utility and is used in guidance provided in the following recommendations.

`UFW` is a user-friendly frontend firewall utility. It uses configurable default policies, and allows administrators to create specific rules to permit or deny connections based on applications, ports, or IP addresses. You can enable the firewall, view its status, and manage rules using simple command-line tools to secure your system.

Warning: 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.

Internal Only - General

4.1.1 Ensure ufw is installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The Uncomplicated Firewall (ufw) is a frontend for iptables and is particularly well-suited for host-based firewalls. ufw provides a framework for managing netfilter, as well as a command-line interface for manipulating the firewall

Rationale:

UFW acts as a frontend for both iptables and nftables and can use either as its backend, though the specific backend depends on the Linux distribution and system configuration.

You can enable the firewall, view its status, and manage rules using simple command-line tools to secure your system.

Impact:

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

Audit:

Run the following command to verify that Uncomplicated Firewall (UFW) is installed:

```
# dpkg-query -s ufw &>/dev/null && echo "ufw is installed"

ufw is installed
```

Remediation:

Run the following command to install Uncomplicated Firewall (UFW):

```
# apt install ufw
```

References:

1. GRID: NIX-00000182
2. CCI-002314
3. NIST SP 800-53 :: AC-17 (1)
4. STIG ID: UBTU-22-251010 | Rule ID: SV-260514r958672 | CAT II
5. STIG ID: UBTU-24-100300 | Rule ID: SV-270654r1067143 | CAT II

Internal Only - General

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.	●	●	●
v8	4.5 Implement and Manage a Firewall on End-User Devices Implement and manage a host-based firewall or port-filtering tool on end-user devices, with a default-deny rule that drops all traffic except those services and ports that are explicitly allowed.	●	●	●
v7	9.4 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.	●	●	●

Internal Only - General

4.1.2 Ensure ufw service is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The service, `ufw.service`, manages UFW. This service is responsible for applying and maintaining firewall rules, controlling which network traffic is allowed in and out of the system.

Note:

- When running ufw enable or starting `ufw.service`, ufw will flush its chains. This is required so ufw can maintain a consistent state, but it may drop existing connections, eg ssh. ufw does support adding rules before enabling the firewall.
- Once ufw is 'enabled', ufw will not flush the chains when adding or removing rules (but will when modifying a rule or changing the default policy)
- By default ufw will prompt when enabling the firewall while running under ssh. This can be disabled by using `ufw --force enable`

Warning: If a rule for openSSH server does not exist, and openSSH is used to administer the host, the ability to administer and/or connect to the host may be lost.

Example command to create a rule to allow ssh:

```
# ufw allow proto tcp from any to any port 22
```

This rule will allow connection to the ssh server from any location. It is highly recommended that this rule be modified to restrict ssh access to only required hosts and to follow local site policy.

- The rules will still be flushed, but the ssh port will be open after enabling the firewall. Please note that once ufw is 'enabled', ufw will not flush the chains when adding or removing rules (but will when modifying a rule or changing the default policy)
- By default, ufw will prompt when enabling the firewall while running under ssh. This can be disabled by using `ufw --force enable`

Rationale:

The ufw service must be enabled and running in order for ufw to protect the system

Impact:

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

Internal Only - General

Audit:

Run the following command to verify `ufw.service` is enabled:

```
# systemctl is-enabled ufw.service

enabled
```

Run the following command to verify `ufw.service` is active:

```
# systemctl is-active ufw.service

active
```

Run the following command to verify ufw is active

```
# ufw status

Status: active
```

Remediation:

Run the following command to unmask the `ufw.service`:

```
# systemctl unmask ufw.service
```

Run the following command to enable and start the `ufw.service`:

```
# systemctl --now enable ufw.service

active
```

Run the following command to enable `ufw`:

```
# ufw enable
```

Default Value:

Inactive

References:

1. GRID: NIX-00000183
2. CCI-002314
3. NIST SP 800-53 :: AC-17 (1)
4. http://manpages.ubuntu.com/manpages/precise/en/man8/ufw.8.html
5. STIG ID: UBTU-20-010434 | Rule ID: SV-238355r958672 | CAT II
6. STIG ID: UBTU-22-251015 | Rule ID: SV-260515r958672 | CAT II
7. STIG ID: UBTU-24-100310 | Rule ID: SV-270655r1067145 | CAT II

Internal Only - General

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.	●	●	●
v8	4.5 Implement and Manage a Firewall on End-User Devices Implement and manage a host-based firewall or port-filtering tool on end-user devices, with a default-deny rule that drops all traffic except those services and ports that are explicitly allowed.	●	●	●
v7	9.4 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.	●	●	●

Internal Only - General

4.1.3 Ensure ufw incoming default is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The default policy for incoming traffic determines if connections attempting to reach your server from external sources will be allowed without requiring specific `UFW` rules.

Rationale:

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

Impact:

Any port or protocol without a explicit allow before the default deny will be blocked. A deny policy silently drops packs causing an eventual time out with no response while a 'reject' policy sends an immediate and explicit "connection refused" or "destination unreachable" message in response.

The following is and example of a command to create a rule to allow connection to the ssh server on the host. This rule should be considered before applying the default deny.

Example:

```
# ufw allow ssh
```

Audit:

Run the following command to verify that the default policy for `incoming` is `deny` or `reject`:

```
# ufw status verbose | awk -F',' '$1~/Default/ {print $1}'
```

Example output:

```
Default: deny (incoming)
```

Remediation:

Run the following command to set the default for `incoming` to `deny`:

```
# ufw default deny incoming
```

Warning: Any port or protocol without an explicit allow before the default deny will be blocked.

Internal Only - General

References:

1. GRID: NIX-00000184
2. NIST SP 800-53 Rev. 5: SC-7

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.4 Implement and Manage a Firewall on Servers Implement and manage a firewall on servers, where supported. Example implementations include a virtual firewall, operating system firewall, or a third-party firewall agent.	●	●	●
v8	4.5 Implement and Manage a Firewall on End-User Devices Implement and manage a host-based firewall or port-filtering tool on end-user devices, with a default-deny rule that drops all traffic except those services and ports that are explicitly allowed.	●	●	●
v7	9.4 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.	●	●	●

Internal Only - General

4.1.4 Ensure ufw outgoing default is configured (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

The default policy for outgoing traffic determines if applications and services running on your server can initiate connections to external networks without requiring specific `UFW` rules.

Rationale:

With a default deny outgoing policy the firewall will deny **ALL** outgoing traffic and is a highly restrictive policy that requires the addition of specific allow rules.

Impact:

Any port and protocol not explicitly allowed will be blocked. The following rules are an example of some outgoing allow rules that should be considered before applying this default deny.

```
ufw allow out http
ufw allow out https
ufw allow out ntp # Network Time Protocol
ufw allow out to any port 53 # DNS
ufw allow out to any port 853 # DNS over TLS
ufw logging on
```

Audit:

Run the following command and verify that the default policy for `outgoing` is `deny` or `reject`:

```
# ufw status verbose | awk -F',' '$1~/Default/ {print $2}'
```

Example output:

```
 deny (outgoing)
```

Remediation:

Run the following command to set the default for `outgoing` to `deny`:

```
# ufw default deny outgoing
```

Warning: Any port or protocol without an explicit allow before the default deny will be blocked.

Internal Only - General

Default Value:

allow (outgoing)

References:

1. GRID: NIX-00000185
2. CCI-000382
3. NIST SP 800-53 :: CM-7
4. NIST SP 800-53A :: CM-7.1 (iii)

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.	●	●	●
v8	4.5 Implement and Manage a Firewall on End-User Devices Implement and manage a host-based firewall or port-filtering tool on end-user devices, with a default-deny rule that drops all traffic except those services and ports that are explicitly allowed.	●	●	●
v7	9.4 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.	●	●	●

Internal Only - General

4.1.5 Ensure ufw routed default is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The default policy for routed traffic determines if `UFW` forwards traffic between different network interfaces without requiring specific UFW rules.

Rationale:

A default `disabled` or `deny` policy ensures that UFW does not forward traffic between different network interfaces by default. This reduces the risk from unwanted or malicious routed traffic.

Impact:

Any port and protocol will be prevented for being routed.

Audit:

Run the following command and verify that the default policy for `routed` is `deny`:

```
# ufw status verbose | awk '/^Default:/{print}' | grep -w routed
```

Example output:

```
 deny (routed)
```

Note: If `net.ipv4.ip_forward` is disabled at the kernel level, `ufw status verbose` may display `disabled (routed)` instead of `deny (routed)`. This reflects the kernel forwarding state, not the UFW routed policy setting. If forwarding is disabled, run `grep -i routed /etc/default/ufw` to verify the UFW policy directly.

Important: Any port and protocol will be prevented from being routed.

Remediation:

Run the following command to set the default for `routed` to `deny`:

```
# ufw default deny routed
```

Default Value:

disabled

References:

1. GRID: NIX-00000186

Internal Only - General

2. CCI-000382
3. NIST SP 800-53 :: CM-7
4. NIST SP 800-53A :: CM-7.1 (iii)

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.	●	●	●
v8	4.5 Implement and Manage a Firewall on End-User Devices Implement and manage a host-based firewall or port-filtering tool on end-user devices, with a default-deny rule that drops all traffic except those services and ports that are explicitly allowed.	●	●	●
v7	9.4 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.	●	●	●

5 Access Control

Internal Only - General

5.1 Configure SSH Server

> **Note:** The recommendations in this section only apply if SSH server is installed on the system. If remote access is not required the SSH server can be removed and this section skipped.

Secure Shell (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.

`sshd_config`:

- First occurrence of an option takes precedence.
- The default `Include` location is the `/etc/ssh/sshd_config.d` directory.
- The audits of the running configuration in this section are run in the context of the root user, the local host name, and the local host's IP address. If a `Match` block exists that matches one of these criteria, the output of the audit will be from the match block. The respective matched criteria should be replaced with a non-matching substitution.
- Defaults:
  - The file `/etc/ssh/sshd_config` includes the line: `Include /etc/ssh/sshd_config.d/*.conf`. This line must appear before any lines containing the `Cipher` argument
  - This directory `/etc/ssh/sshd_config.d/` includes a file `/etc/ssh/sshd_config.d/50-redhat.conf`
  - The file `/etc/ssh/sshd_config.d/50-redhat.conf` includes the line `Include /etc/crypto-policies/back-ends/opensshserver.config`
- `Ciphers`, `MACs`, and `KexAlgorithms`:
  - Though these options may be configured through the `Ciphers`, `MACs`, and `KexAlgorithms` options in the `/etc/ssh/sshd_config` file, it is **highly** recommended that they are configured through system-wide-crypto-policy.
  - If the recommendations in the subsection "Configure system wide crypto policy" have been followed, the audits should be in a passing state. Please review that section before following the Remediation Procedure for these Recommendations.
  - By default, system-wide-crypto-policy is applied to the openSSH server. If the following defaults don't exist due to modifications or upgrade from a earlier release, the system-wide-crypto-policy may not be included by the openSSH server. It is recommended that these defaults be restored, created, or the line `Include /etc/crypto-policies/back-ends/opensshserver.config` be added before any lines containing the `Cipher`, `MACs`, or `KexAlgorithms` arguments.
- `Include`:
  - Include the specified configuration file(s).
  - Multiple pathnames may be specified and each pathname may contain glob(7) wildcards that will be expanded and processed in lexical order.

Internal Only - General

    - Files without absolute paths are assumed to be in `/etc/ssh/`.
    - An Include directive may appear inside a Match block to perform conditional inclusion.
- `Match`:
    - Introduces a conditional block. If all of the criteria on the Match line are satisfied, the keywords on the following lines override those set in the global section of the config file, until either another Match line or the end of the file. If a keyword appears in multiple Match blocks that are satisfied, only the first instance of the keyword is applied.
    - The arguments to Match are one or more criteria-pattern pairs or the single token All which matches all criteria. The available criteria are User, Group, Host, LocalAddress, LocalPort, and Address.
    - The match patterns may consist of single entries or comma-separated lists and may use the wildcard and negation operators described in the PATTERNS section of ssh_config(5).
    - The patterns in an Address criteria may additionally contain addresses to match in CIDR address/masklen format, such as `192.0.2.0/24` or `2001:db8::/32`. Note that the mask length provided must be consistent with the address - it is an error to specify a mask length that is too long for the address or one with bits set in this host portion of the address. For example, `192.0.2.0/33` and `192.0.2.0/8`, respectively.
    - Only a subset of keywords may be used on the lines following a Match keyword. Available keywords are available in the ssh_config man page.
- Once changes have been made to the `sshd` configuration, the `sshd.service` must be reloaded.

Command to re-load the SSH daemon configuration:

```
# systemctl reload-or-restart sshd.service
```

`sshd` command:

- `-T` - Extended test mode. Check the validity of the configuration file, output the effective configuration to stdout and then exit. Optionally, Match rules may be applied by specifying the connection parameters using one or more `-C` options.
- `-C` - connection_spec. Specify the connection parameters to use for the -T extended test mode. If provided, any Match directives in the configuration file that would apply are applied before the configuration is written to standard output. The connection parameters are supplied as keyword=value pairs and may be supplied in any order, either with multiple `-C` options or as a comma-separated list. The keywords are `addr`, `user`, `host`, `laddr`, `lport`, and `rdomain` and correspond to source address, user, resolved source host name, local address, local port number and routing domain respectively.

Internal Only - General

5.1.1 Ensure access to /etc/ssh/sshd_config is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`/etc/ssh/sshd_config` is the main configuration file for the OpenSSH server daemon (sshd), which manages secure remote logins and other SSH services.

`/etc/ssh/sshd_config` is the default location for drop-in configuration files used by sshd.

More information about the openSSH server configuration is available in the "Configure SSH Server" section overview.

Rationale:

Configuration specifications for `sshd` need to be protected from unauthorized changes by non-privileged users.

Audit:

Run the following command to verify access to `/etc/ssh/sshd_config` is configured:

```
# stat -Lc 'File: "%n" Mode: "%#a"  Owner: "%U" Group: "%G"'
/etc/ssh/sshd_config
```

Verify mode is `0600` or more restrictive, owner is `root`, and group is `root`.

Example output:

```
File: "/etc/ssh/sshd_config" Mode: "0600"  Owner: "root" Group: "root"
```

Run the following script to verify files in `/etc/ssh/sshd_config.d/` ending in `.conf` are configured:

- Mode `0600` or more restrictive
- Owned by the `root` user
- Group owned by the group `root`.

Internal Only - General

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

{
   while IFS= read -r -d $'\0' l_file; do
      stat -Lc 'File: "%n" Mode: "%#a"  Owner: "%U" Group: "%G"' "$l_file"
   done < <(find /etc/ssh/sshd_config.d/ -xdev -mount -type f -name '*.conf'
\
   \( -perm /077 -o ! -user root -o ! -group root \) -print0 2>/dev/null)
}
```

Nothing should be returned

> **Note:** If other locations are listed in an `Include` statement, `*.conf` files in these locations should also be checked.

Remediation:

Run the following script to set ownership and permissions on `/etc/ssh/sshd_config` and files ending in `.conf` in the `/etc/ssh/sshd_config.d` directory:

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

{
   chmod u-x,og-rwx /etc/ssh/sshd_config
   chown root:root /etc/ssh/sshd_config
   while IFS= read -r -d $'\0' l_file; do
      if [ -e "$l_file" ]; then
         chmod u-x,og-rwx "$l_file"
         chown root:root "$l_file"
      fi
   done < <(find /etc/ssh/sshd_config.d -type f -print0 2>/dev/null)
}
```

> **Note:** If other locations are listed in an `Include` statement, `*.conf` files in these locations access should also be modified.

Default Value:

Access: (0600/-rw-------) Uid: (0/ root) Gid: (0/ root)

References:

1. GRID: NIX-00000187
2. NIST SP 800-53 Rev. 5: AC-3. MP-2

Internal Only - General

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.	●	●	●

Internal Only - General

5.1.2 Ensure access to /etc/ssh/sshd_config.d/ directory is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`/etc/ssh/sshd_config.d/` is the default drop-in configuration directory for the OpenSSH server daemon (`sshd`). The line `Include /etc/ssh/sshd_config.d/*.conf` shipped in the default `/etc/ssh/sshd_config` causes every file in this directory ending in `.conf` to be processed as part of the effective server configuration, with first-occurrence-wins precedence over directives that follow the `Include` in the parent file.

This recommendation governs the **directory itself**: ownership of `/etc/ssh/sshd_config.d/`, and the mode bits applied to the directory inode. It complements recommendation 5.1.1 (`Ensure access to /etc/ssh/sshd_config is configured`), which governs the parent file and the `*.conf` drop-in files **inside** this directory but does not constrain the directory entry.

More information about the openSSH server configuration is available in the "Configure SSH Server" section overview.

Rationale:

A drop-in configuration directory at the default mode of `0755 root:root` is world-readable and world-traversable. Any local user can therefore:

1. Enumerate the per-host SSH server policy
2. Read the contents of any `.conf` drop-in that itself is world-readable
3. Reconnoitre for related artifacts: AllowGroups lists, ListenAddress restrictions, banner paths, alternative `HostKey` paths, and `AuthorizedKeysFile` overrides each reveal information useful to an attacker performing local reconnaissance prior to privilege escalation or lateral movement.

Configuration specifications for `sshd` — including the per-host overrides contained in this directory — need to be protected from unauthorized disclosure as well as unauthorized modification by non-privileged users. Restricting the directory to mode `0700 root:root` ensures that only the `root` user can list, traverse, or modify the directory, matching the protection already required for the parent `sshd_config` file and the `.conf` files inside the directory.

Internal Only - General

Audit:

Run the following command to verify ownership and mode on `/etc/ssh/sshd_config.d/` are configured:

```
# find /etc/ssh/sshd_config.d/ -maxdepth 0 \( ! -user root -o ! -group root -
o -perm /077 \)
```

Nothing should be returned.

> **Note:** If the SSH server has been configured to read drop-in files from an `Include` location other than `/etc/ssh/sshd_config.d/`, the additional directory or directories listed by that `Include` statement should also be checked against the same criteria.

Example identification of additional `Include` locations:

```
# sshd -T 2>/dev/null | awk '$1 == "include" {print $2}'
```

For each path returned, repeat the audit substituting the path for `/etc/ssh/sshd_config.d/` and adjusting `-maxdepth` if the path is a glob or contains a wildcard.

Remediation:

Run the following commands to set ownership and mode on `/etc/ssh/sshd_config.d/`:

```
# chown root:root /etc/ssh/sshd_config.d/
# chmod 0700 /etc/ssh/sshd_config.d/
```

> **Note:** If the SSH server reads drop-in files from an additional `Include` location, apply the same ownership and mode to each such directory.

No service reload is required — the directory mode does not affect the running `sshd` process, only future reads of the directory.

References:

1. GRID: NIX-00000408
2. NIST SP 800-53 Rev. 5: AC-3, MP-2
3. https://manpages.debian.org/trixie/openssh-server/sshd_config.5.en.html
4. https://manpages.debian.org/trixie/openssh-server/sshd.8.en.html

Internal Only - General

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.	●	●	●

Internal Only - General

5.1.3 Ensure access to SSH private host key files is 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 threat actor obtains the private SSH host key file, the host could be impersonated.

Impact:

If the private keys used by the openSSH server have the incorrect mode, owner, or group owner the server service may not start.

Audit:

Verify SSH private host key files are:

- Mode `0600`.
- Owned by the `root` user.
- Group owned by the group `root`.

Run the following script to generate a list of SSH private host key files:

Internal Only - General

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

{
   l_sshd_cmd="$(readlink -e /usr/sbin/sshd || readlink -e /sbin/sshd)"
   l_keygen="$(readlink -e /usr/bin/ssh-keygen || readlink -e /bin/ssh-
keygen)"
   l_sshd_config="$(readlink -e /etc/ssh/sshd_config)"
   l_sshd_args=(-T)
   grep -qiE '^\s*Match\s+(User|Group|Host|LocalAddress|LocalPort|Address)' \
      "$l_sshd_config" 2>/dev/null && l_sshd_args+=(-C
"user=root,host=localhost,addr=127.0.0.1")
   mapfile -t l_hostkeys < <("$l_sshd_cmd" "${l_sshd_args[@]}" 2>/dev/null |
awk '$1=="hostkey"{print $2}')
   for l_file in "${l_hostkeys[@]}"; do
      "$l_keygen" -lf "$l_file" &>/dev/null && \
         stat -Lc 'File: "%n" Mode: "%#a" Owner: "%U" Group: "%G"' "$l_file"
   done
}
```

Example output:

```
File: "/etc/ssh/ssh_host_rsa_key" Mode: "0600"  Owner: "root" Group: "root"
File: "/etc/ssh/ssh_host_ecdsa_key" Mode: "0600"  Owner: "root" Group: "root"
File: "/etc/ssh/ssh_host_ed25519_key" Mode: "0600"  Owner: "root" Group:
"root"
```

Verify all files in the output are:

- Mode `0600`.
- Owner by the user `root`.
- Group owned by the group `root`.

Remediation:

Update the access to the private keys being used by the open SSH server.

- mode to `0600`.
- owner to the user `root`.
- group owner to the group `root`.

Run the following script to update the access on the private keys used by the open SSH server:

Internal Only - General

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

{
   l_sshd_cmd="$(readlink -e /usr/sbin/sshd || readlink -e /sbin/sshd)"
   l_keygen="$(readlink -e /usr/bin/ssh-keygen || readlink -e /bin/ssh-
keygen)"
   while IFS= read -r l_file; do
      if "$l_keygen" -lf &>/dev/null "$l_file"; then
         chown root:root "$l_file"
         chmod 0600 "$l_file"
      fi
   done < <("$l_sshd_cmd" -T | awk '$1=="hostkey" {print $2".pub"}'
2>/dev/null)
}
```

Default Value:

Mode: "0600" Owner: "root" Group: "root"

References:

1. GRID: NIX-00000189
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-08-010490 | RULE ID: SV-230287r1017098 | CAT II
5. STIG ID: RHEL-09-255120 | RULE ID: SV-258000r1045063 | CAT II
6. STIG ID: ALMA-09-020920 | RULE ID: SV-269265r1050147 | CAT II
7. STIG ID: SLES-15-040250 | RULE ID: SV-235009r991589 | CAT II
8. CCI-000366

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.	●	●	●

Internal Only - General

5.1.4 Ensure access to SSH public host key files is 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:

Verify SSH public host key files are:

- Mode `0644` or more restrictive.
- Owned by the `root` user.
- Group owned by the group `root`.

Run the following script to generate a list of SSH public host key files:

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

{
   l_sshd_cmd="$(readlink -e /usr/sbin/sshd || readlink -e /sbin/sshd)"
   l_keygen="$(readlink -e /usr/bin/ssh-keygen || readlink -e /bin/ssh-
keygen)"
   l_sshd_config="$(readlink -e /etc/ssh/sshd_config)"
   l_sshd_args=(-T)
   grep -qiE '^\s*Match\s+(User|Group|Host|LocalAddress|LocalPort|Address)' \
      "$l_sshd_config" 2>/dev/null && l_sshd_args+=(-C
"user=root,host=localhost,addr=127.0.0.1")
   mapfile -t l_hostkeys < <("$l_sshd_cmd" "${l_sshd_args[@]}" 2>/dev/null |
awk '$1=="hostkey" {print $2".pub"}' 2>/dev/null)
   for l_file in "${l_hostkeys[@]}"; do
      "$l_keygen" -lf "$l_file" &>/dev/null && \
         stat -Lc 'File: "%n" Mode: "%#a"  Owner: "%U" Group: "%G"' "$l_file"
   done
}
```

Example output:

Internal Only - General

```
File: "/etc/ssh/ssh_host_rsa_key.pub" Mode: "0644"  Owner: "root" Group:
"root"
File: "/etc/ssh/ssh_host_ecdsa_key.pub" Mode: "0644"  Owner: "root" Group:
"root"
File: "/etc/ssh/ssh_host_ed25519_key.pub" Mode: "0644"  Owner: "root" Group:
"root"
```

Verify all files in the output are:

- Mode `0644`.
- Owner by the user `root`.
- Group owned by the group `root`.

Remediation:

Update the access to the public keys being used by the open SSH server.

- mode to `0644`.
- owner to the user `root`.
- group owner to the group `root`.

Run the following script to update the access on the public keys used by the open SSH server:

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

{
   l_sshd_cmd="$(readlink -e /usr/sbin/sshd || readlink -e /sbin/sshd)"
   l_keygen="$(readlink -e /usr/bin/ssh-keygen || readlink -e /bin/ssh-
keygen)"
   while IFS= read -r l_file; do
      if "$l_keygen" -lf &>/dev/null "$l_file"; then
         chown root:root "$l_file"
         chmod 0644 "$l_file"
      fi
   done < <("$l_sshd_cmd" -T | awk '$1=="hostkey" {print $2".pub"}'
2>/dev/null)
}
```

Default Value:

Mode: "0644" Owner: "root" Group: "root"

References:

1. GRID: NIX-00000190
2. NIST SP 800-53 Rev. 5: AC-3. MP-2

Internal Only - General

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.	●	●	●

Internal Only - General

5.1.5 Ensure sshd access is configured (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.

More information about the openSSH server configuration is available in the "Configure SSH Server" section overview.

Rationale:

Restricting which users can remotely access the system via SSH will help ensure that only authorized users access the system.

Internal Only - General

Impact:

These options are "ANDed" together. If both `AllowUsers` and `AllowGroups` are set, connections will be limited to the list of users that are also a member of an allowed group. It is recommended that only one be set for clarity and ease of administration.

Audit:

Run the following command and verify the output:

```
# sshd -T | grep -Pi -- '^\h*(allow|deny)(users|groups)\h+\H+'
```

Verify that the output matches at least one of the following lines:

```
allowusers <userlist>
-OR-
allowgroups <grouplist>
-OR-
denyusers <userlist>
-OR-
denygroups <grouplist>
```

Review the list(s) to ensure included users and/or groups follow local site policy

- IF - `Match` set statements are used in your environment, specify the connection parameters to use for the `-T` extended test mode and run the audit to verify the setting is not incorrectly configured in a match block

Example additional audit needed for a match block for the user `sshuser`:

```
# sshd -T -C user=sshuser | grep -Pi --
'^\h*(allow|deny)(users|groups)\h+\H+'
```

Remediation:

Create or edit a `*.conf` file in the `/etc/ssh/sshd_config.d/` and set one or more of the parameters above any `Include` and `Match` set statements as follows:

```
AllowUsers <userlist>
 - AND/OR -
AllowGroups <grouplist>
```

> **Note:** It is easier to manage an allow list than a deny list. In a deny list, you could potentially add a user or group and forget to add it to the deny list.

Run the following command to apply the configuration change to the running sshd daemon without interrupting active sessions:

```
# systemctl reload sshd
```

Default Value:

None

Internal Only - General

References:

1. GRID: NIX-00000191
2. NIST SP 800-53 Rev. 5: AC-3. MP-2
3. SSHD(8)
4. https://documentation.suse.com/en-us/sles/15-SP6/html/SLES-all/cha-ssh.html
5. SSHD_CONFIG(5)

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	4.3 Ensure the Use of Dedicated Administrative Accounts 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.	●	●	●

Internal Only - General

5.1.6 Ensure sshd 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.

More information about the openSSH server configuration is available in the "Configure SSH Server" section overview.

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.

Not having a properly configured SSH banner, or leaving it at the default value, can expose an organization to legal and compliance risks by failing to provide users with necessary access policies or warnings about unauthorized access. It may violate regulations requiring user acknowledgment of system usage rules, reduce security awareness, and leave systems vulnerable to malicious or unintended access.

Audit:

1. Ensure the the `Banner` option is set.

Run the following command to verify `Banner` is set:

```
# sshd -T | grep -Pi -- '^banner\h+\/\H+'
```

Example:

```
banner /etc/issue.net
```

- IF - `Match` set statements are used in your environment, specify the connection parameters to use for the `-T` extended test mode and run the audit to verify the setting is not incorrectly configured in a match block

Example additional audit needed for a match block for the user `sshuser`:

```
# sshd -T -C user=sshuser | grep -Pi -- '^banner\h+\/\H+'
```

2. Ensure the banner matches local site policy.

Run the following command and verify that the contents or the file being called by the `Banner` argument match site policy:

Internal Only - General

```
# [ -e "$(sshd -T | awk '$1 == "banner" {print $2}')" ] && cat "$(sshd -T |
awk '$1 == "banner" {print $2}')"
```

Run the following command and verify no results are returned:

```
# grep -Psi -- "(\\\v|\\\r|\\\m|\\\s|\b$(grep '^ID=' /etc/os-release | cut -
d= -f2 | sed -e 's/"//g')\b)" "$(sshd -T | awk '$1 == "banner" {print $2}')"
```

Remediation:

Create or edit a `*.conf` file in the `/etc/ssh/sshd_config.d/` and set the `Banner` parameter above any `Include` and `Match` entries as follows:

```
Banner /etc/issue.net
```

Edit the file being called by the `Banner` argument with the appropriate contents according to your site policy, remove any instances of `\m` , `\r` , `\s` , `\v` or references to the `OS platform`

Example:

```
# printf '%s\n' "Authorized users only. All activity may be monitored and
reported." > "$(sshd -T | awk '$1 == "banner" {print $2}')"
```

Run the following command to apply the configuration change to the running sshd daemon without interrupting active sessions:

```
# systemctl reload sshd
```

References:

1. GRID: NIX-00000192
2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

5.1.7 Ensure sshd Ciphers are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Secure Shell (SSH) ciphers are encryption algorithms used by the SSH protocol to secure data transmitted between a client and a server. They protect the privacy and integrity of the connection, using symmetric encryption to encrypt data after a secure session is established. During connection setup, the client and server negotiate to select the strongest available cipher that both support, ensuring the connection is as secure as possible

More information about the openSSH server configuration is available in the "Configure SSH Server" section overview.

Rationale:

Weak ciphers that are used for authentication to the cryptographic module cannot be relied upon to provide confidentiality or integrity, and system data may be compromised.

Impact:

A too restrictive list of ciphers can lead to the "no matching cipher found" error during an SSH connection. This error indicates that the SSH client and server are unable to agree on a common cipher to use for the secure communication. This typically occurs due to a mismatch in supported ciphers or when older, less secure ciphers are disabled on one side but still expected by the other.

Audit:

Run the following command to verify none of the "weak" ciphers are being used:

```
# sshd -T | grep -Pi --
'^ciphers\h+\"?([^#\n\r]+,)?((3des|blowfish|cast128|aes(128|192|256))-
cbc|arcfour(128|256)?|rijndael-cbc@lysator\.liu\.se)\b'
```

Nothing should be returned.

Run the following command to determine if the `chacha20-poly1305@openssh.com` cipher is being used.

```
#  sshd -T 2>&1 | awk '($1=="ciphers" && $2~/chacha20-
poly1305@openssh\.com/){print}'
```

If any line is returned, review `CVE-2023-48795` and verify the system has been patched.

Internal Only - General

Remediation:

Edit or create a drop-in file under `/etc/ssh/sshd_config.d/` to remove weak ciphers using the subtractive form. The drop-in file should be named so that lexical order places it before any other drop-in that sets `Ciphers`:

```
Ciphers -3des-cbc,blowfish-cbc,cast128-cbc,aes128-cbc,aes192-cbc,aes256-
cbc,arcfour,arcfour128,arcfour256,rijndael-cbc@lysator.liu.se
```

Example `10-cis-sshd.conf`:

```
# printf '%s\n' 'Ciphers -3des-cbc,blowfish-cbc,cast128-cbc,aes128-
cbc,aes192-cbc,aes256-cbc,arcfour,arcfour128,arcfour256,rijndael-
cbc@lysator.liu.se' >> /etc/ssh/sshd_config.d/10-cis-sshd.conf
```

Apply the change to the running daemon:

```
# systemctl reload-or-restart sshd.service
```

Note:

- The list above contains the cipher algorithms that this recommendation prohibits. If the site has additional ciphers it wishes to disable (e.g., to align with a site-specific crypto policy), append them to the `Ciphers -` argument, comma-separated, with no spaces. The subtractive form removes the listed algorithms from OpenSSH's compiled-in default set, so future additions to that default set are inherited rather than frozen.
- Verify the effective post-include configuration with `sshd -T | grep -i ciphers`. The `sshd -T` output is the authoritative source for what the running daemon will negotiate.

Default Value:

Ciphers aes256-gcm@openssh.com,chacha20-poly1305@openssh.com,aes256-ctr,aes128-gcm@openssh.com,aes128-ctr

References:

1. GRID: NIX-00000193
2. https://nvd.nist.gov/vuln/detail/CVE-2019-1543
3. https://nvd.nist.gov/vuln/detail/CVE-2016-2183
4. https://nvd.nist.gov/vuln/detail/CVE-2008-5161
5. https://www.openssh.com/txt/cbc.adv
6. https://www.openssh.com/txt/cbc.adv
7. SSHD_CONFIG(5)
8. NIST SP 800-53 Rev. 5: SC-8
9. https://nvd.nist.gov/vuln/detail/CVE-2023-48795

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.10 Encrypt Sensitive Data in Transit Encrypt sensitive data in transit. Example implementations can include: Transport Layer Security (TLS) and Open Secure Shell (OpenSSH).		●	●
v7	14.4 Encrypt All Sensitive Information in Transit Encrypt all sensitive information in transit.		●	●

Internal Only - General

5.1.8 Ensure sshd ClientAliveInterval and ClientAliveCountMax are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

> **Note:** To clarify, the two settings described below are only meant for idle connections from a protocol perspective and are not meant to check if the user is active or not. An idle user does not mean an idle connection. SSH does not and never had, intentionally, the capability to drop idle users. In SSH versions before `8.2p1` there was a bug that caused these values to behave in such a manner that they were abused to disconnect idle users. This bug has been resolved in `8.2p1` and thus it can no longer be abused disconnect idle users.

The two options `ClientAliveInterval` and `ClientAliveCountMax` control the timeout of SSH sessions. Taken directly from `man 5 sshd_config`:

- `ClientAliveInterval` Sets a timeout interval in seconds after which if no data has been received from the client, sshd(8) will send a message through the encrypted channel to request a response from the client. The default is 0, indicating that these messages will not be sent to the client.
- `ClientAliveCountMax` Sets the number of client alive messages which may be sent without sshd(8) receiving any messages back from the client. If this threshold is reached while client alive messages are being sent, sshd will disconnect the client, terminating the session. It is important to note that the use of client alive messages is very different from TCPKeepAlive. The client alive messages are sent through the encrypted channel and therefore will not be spoofable. The TCP keepalive option en-abled by TCPKeepAlive is spoofable. The client alive mechanism is valuable when the client or server depend on knowing when a connection has become unresponsive.

  The default value is 3. If ClientAliveInterval is set to 15, and ClientAliveCountMax is left at the default, unresponsive SSH clients will be disconnected after approximately 45 seconds. Setting a zero ClientAliveCountMax disables connection termination.

More information about the openSSH server configuration is available in the "Configure SSH Server" section overview.

Internal Only - General

Rationale:

In order to prevent resource exhaustion, appropriate values should be set for both `ClientAliveInterval` and `ClientAliveCountMax`. Specifically, looking at the source code, `ClientAliveCountMax` must be greater than zero in order to utilize the ability of SSH to drop idle connections. If connections are allowed to stay open indefinitely, this can potentially be used as a DDOS attack or simple resource exhaustion could occur over unreliable networks.

The example set here is a 45 second timeout. Consult your site policy for network timeouts and apply as appropriate.

Audit:

Run the following command and verify `ClientAliveInterval` and `ClientAliveCountMax` are greater than zero:

```
# sshd -T | grep -Pi -- '(clientaliveinterval|clientalivecountmax)'
```

Example Output:

```
clientaliveinterval 15
clientalivecountmax 3
```

- IF - `Match` set statements are used in your environment, specify the connection parameters to use for the `-T` extended test mode and run the audit to verify the setting is not incorrectly configured in a match block

Example additional audit needed for a match block for the user `sshuser`:

```
# sshd -T -C user=sshuser | grep -Pi --
'(clientaliveinterval|clientalivecountmax)'
```

Note: If provided, any Match directives in the configuration file that would apply are applied before the configuration is written to standard output. The connection parameters are supplied as keyword=value pairs and may be supplied in any order, either with multiple `-C` options or as a comma-separated list. The keywords are `addr` (source address), `user` (user), `host` (resolved source host name), `laddr` (local address), `lport` (local port number), and `rdomain` (routing domain).

Remediation:

Create or edit a `*.conf` file in the `/etc/ssh/sshd_config.d/` and set the `ClientAliveInterval` and `ClientAliveCountMax` parameters above any `Include` and `Match` entries according to site policy.

Example:

```
ClientAliveInterval 15
ClientAliveCountMax 3
```

Run the following command to apply the configuration change to the running sshd daemon without interrupting active sessions:

Internal Only - General

```
# systemctl reload sshd
```

Default Value:

ClientAliveInterval 0

ClientAliveCountMax 3

References:

1. GRID: NIX-00000194
2. SSHD(8)
3. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5
4. SSHD_CONFIG(5)

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

5.1.9 Ensure sshd DisableForwarding is enabled (Automated)

Profile Applicability:

- Level 1 - Workstation
- Level 1 - Server

Description:

The `DisableForwarding` parameter disables all forwarding features, including X11, ssh-agent(1), TCP and StreamLocal. This option overrides all other forwarding-related options and may simplify restricted configurations.

- X11Forwarding provides the ability to tunnel X11 traffic through the connection to enable remote graphic connections.
- ssh-agent is a program to hold private keys used for public key authentication. Through use of environment variables the agent can be located and automatically used for authentication when logging in to other machines using ssh.
- 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.

More information about the openSSH server configuration is available in the "Configure SSH Server" section overview.

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.

anyone with root privilege on the the intermediate server can make free use of ssh-agent to authenticate them to other servers

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.

Internal Only - General

Impact:

SSH tunnels are widely used in many corporate environments. In some 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 to verify `DisableForwarding` is set to `yes`:

```
# sshd -T | grep disableforwarding

disableforwarding yes
```

Remediation:

Create or edit a `*.conf` file in the `/etc/ssh/sshd_config.d/` and set the `DisableForwarding` parameter to `yes` above any `Include` entry as follows:

```
DisableForwarding yes
```

Run the following command to apply the configuration change to the running sshd daemon without interrupting active sessions:

```
# systemctl reload sshd
```

References:

1. GRID: NIX-00000196
2. SSHD(8)
3. NIST SP 800-53 Rev. 5: CM-7
4. sshd_config(5)

CIS Controls:

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

Internal Only - General

5.1.10 Ensure sshd GSSAPIAuthentication is disabled (Automated)

Profile Applicability:

- Level 1 - Workstation
- Level 1 - Server

Description:

The `GSSAPIAuthentication` parameter specifies whether user authentication based on GSSAPI is allowed

More information about the openSSH server configuration is available in the "Configure SSH Server" section overview.

Rationale:

Allowing GSSAPI authentication through SSH exposes the system's GSSAPI to remote hosts, and should be disabled to reduce the attack surface of the system

Impact:

Disabling GSSAPI authentication in environments that rely on centralized authentication could disrupt legitimate access.

Audit:

Run the following command to verify `GSSAPIAuthentication` is set to `no`:

```
# sshd -T | grep gssapiauthentication

gssapiauthentication no
```

- IF - `Match` set statements are used in your environment, specify the connection parameters to use for the `-T` extended test mode and run the audit to verify the setting is not incorrectly configured in a match block

Example additional audit needed for a match block for the user `sshuser`:

```
# sshd -T -C user=sshuser | grep gssapiauthentication
```

Remediation:

Create or edit a `*.conf` file in the `/etc/ssh/sshd_config.d/` and set the `GSSAPIAuthentication` parameter to `no` above any `Include` and `Match` entries as follows:

```
GSSAPIAuthentication no
```

Run the following command to apply the configuration change to the running sshd daemon without interrupting active sessions:

Internal Only - General

```
# systemctl reload sshd
```

Default Value:

GSSAPIAuthentication no

References:

1. GRID: NIX-00000197
2. SSHD(8)
3. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5
4. SSHD_CONFIG(5)

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.2 Use Unique Passwords Use unique passwords for all enterprise assets. Best practice implementation includes, at a minimum, an 8-character password for accounts using MFA and a 14-character password for accounts not using MFA.	●	●	●
v7	4.4 Use Unique Passwords Where multi-factor authentication is not supported (such as local administrator, root, or service accounts), accounts will use passwords that are unique to that system.		●	●

Internal Only - General

5.1.11 Ensure sshd 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.

More information about the openSSH server configuration is available in the "Configure SSH Server" section overview.

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 to verify `HostbasedAuthentication` is set to `no`:

```
# sshd -T | grep hostbasedauthentication

hostbasedauthentication no
```

- IF - `Match` set statements are used in your environment, specify the connection parameters to use for the `-T` extended test mode and run the audit to verify the setting is not incorrectly configured in a match block

Example additional audit needed for a match block for the user `sshuser`:

```
# sshd -T -C user=sshuser | grep hostbasedauthentication
```

Remediation:

Create or edit a `*.conf` file in the `/etc/ssh/sshd_config.d/` and set the `HostbasedAuthentication` parameter to `no` above any `Include` and `Match` entries as follows:

```
HostbasedAuthentication no
```

Run the following command to apply the configuration change to the running sshd daemon without interrupting active sessions:

Internal Only - General

```
# systemctl reload sshd
```

Default Value:

HostbasedAuthentication no

References:

1. GRID: NIX-00000198
2. SSHD(8)
3. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5
4. SSHD_CONFIG(5)

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.2 Use Unique Passwords Use unique passwords for all enterprise assets. Best practice implementation includes, at a minimum, an 8-character password for accounts using MFA and a 14-character password for accounts not using MFA.	●	●	●
v7	4.4 Use Unique Passwords Where multi-factor authentication is not supported (such as local administrator, root, or service accounts), accounts will use passwords that are unique to that system.		●	●

Internal Only - General

5.1.12 Ensure sshd 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`.

More information about the openSSH server configuration is available in the "Configure SSH Server" section overview.

Rationale:

Setting this parameter forces users to enter a password when authenticating with SSH.

Audit:

Run the following command to verify `IgnoreRhosts` is set to `yes`:

```
# sshd -T | grep ignorerhosts

ignorerhosts yes
```

Remediation:

Create or edit a `*.conf` file in the `/etc/ssh/sshd_config.d/` and set the `IgnoreRhosts` parameter to `yes` above any `Include` entry as follows:

```
IgnoreRhosts yes
```

Run the following command to apply the configuration change to the running sshd daemon without interrupting active sessions:

```
# systemctl reload sshd
```

Default Value:

IgnoreRhosts yes

References:

1. GRID: NIX-00000199
2. SSHD(8)
3. NIST SP 800-53 Rev. 5: CM-1,CM-2, CM-6, CM-7, IA-5
4. SSHD_CONFIG(5)

Internal Only - General

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.		●	●

Internal Only - General

5.1.13 Ensure sshd KexAlgorithms is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Key exchange is any method in cryptography by which cryptographic keys are exchanged between two parties, allowing use of a cryptographic algorithm. If the sender and receiver wish to exchange encrypted messages, each must be equipped to encrypt messages to be sent and decrypt messages received

More information about the openSSH server configuration is available in the "Configure SSH Server" section overview.

Rationale:

Key exchange methods that are considered weak should be removed. A key exchange method may be weak because too few bits are used, or the hashing algorithm is considered too weak. Using weak algorithms could expose connections to man-in-the-middle attacks

Audit:

Run the following command to verify none of the "weak" Key Exchange algorithms are being used:

```
# sshd -T | grep -Pi -- 'kexalgorithms\h+([^#\n\r]+,)?(diffie-hellman-group1-sha1|diffie-hellman-group14-sha1|diffie-hellman-group-exchange-sha1)\b'
```

Nothing should be returned.

> **Note:** If anything is returned, review the list of Key Exchange Algorithms. The following are considered "weak" Key Exchange Algorithms, and should not be used:

```
diffie-hellman-group1-sha1
diffie-hellman-group14-sha1
diffie-hellman-group-exchange-sha1
```

Remediation:

Create or edit a `*.conf` file in the `/etc/ssh/sshd_config.d/` and add/modify the `KexAlgorithms` line to contain a comma separated list of the site unapproved (weak) KexAlgorithms preceded with a `-` above any `Include` entries:

Example:

Internal Only - General

```
KexAlgorithms -diffie-hellman-group1-sha1,diffie-hellman-group14-sha1,diffie-
hellman-group-exchange-sha1
```

> **Note:** First occurrence of an option takes precedence. If Include locations are enabled, used, and order of precedence is understood in your environment, the entry may be created in a file in Include location.

Run the following command to apply the configuration change to the running sshd daemon without interrupting active sessions:

```
# systemctl reload sshd
```

Default Value:

Version 9.0 - 9.8: KexAlgorithms sntrup761x25519-sha512,sntrup761x25519-sha512@openssh.com,curve25519-sha256,curve25519-sha256@libssh.org,ecdh-sha2-nistp256,ecdh-sha2-nistp384,ecdh-sha2-nistp521

Version 9.9 and greater: KexAlgorithms mlkem768x25519-sha256,sntrup761x25519-sha512,sntrup761x25519-sha512@openssh.com,curve25519-sha256,curve25519-sha256@libssh.org,ecdh-sha2-nistp256,ecdh-sha2-nistp384,ecdh-sha2-nistp521

References:

1. GRID: NIX-00000200
2. NIST SP 800-53 Rev. 5: SC-8
3. SSHD(8)
4. SSHD_CONFIG(5)
5. https://ubuntu.com/server/docs/openssh-crypto-configuration

Additional Information:

The supported algorithms are:

Internal Only - General

```
diffie-hellman-group1-sha1
diffie-hellman-group14-sha1
diffie-hellman-group14-sha256
diffie-hellman-group16-sha512
diffie-hellman-group18-sha512
diffie-hellman-group-exchange-sha1
diffie-hellman-group-exchange-sha256
ecdh-sha2-nistp256
ecdh-sha2-nistp384
ecdh-sha2-nistp521
curve25519-sha256
curve25519-sha256@libssh.org
sntrup761x25519-sha512
sntrup761x25519-sha512@openssh.com
mlkem768x25519-sha256 # Only supported by version 9.9 and greater
```

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.10 Encrypt Sensitive Data in Transit Encrypt sensitive data in transit. Example implementations can include: Transport Layer Security (TLS) and Open Secure Shell (OpenSSH).		●	●
v7	14.4 Encrypt All Sensitive Information in Transit Encrypt all sensitive information in transit.		●	●

Internal Only - General

5.1.14 Ensure sshd post-quantum cryptography key exchange algorithms are configured (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

OpenSSH has integrated post-quantum cryptography (PQC) to secure SSH connections against potential attacks from future quantum computers. This includes the adoption of hybrid key agreement algorithms that combine classical cryptographic methods with quantum-resistant ones.

More information about the openSSH server configuration is available in the "Configure SSH Server" section overview.

Rationale:

PQC helps secure SSH connections against potential attacks from future quantum computers providing protection against both classical and quantum adversaries.

The combination of established classical algorithms with the forward-looking security of PQC algorithms, offers a robust solution during the transition phase to fully quantum-resistant cryptography.

Impact:

A restrictive list of key exchange algorithms available to the SSH server may prevent some clients from being able to connect.

Including a non-supported key exchange algorithm in the `kexalgorithms` option, may prevent the SSH server from starting.

> **Note:** The list of supported key exchange algorithms may also be obtained using the `ssh -Q KexAlgorithms` command.

Audit:

1. Run the following command to determine the version of openSSH-server:

```
# "$(readlink -e /usr/sbin/sshd || readlink -e /sbin/sshd)" -V 2>&1 | \
grep -Psio 'openssh_\d+\.\d+' | awk -F'_' '{print $2}'
```

2. Verify post-quantum cryptography kexalgorithm(s) are available:

Run the following command to verify the `sntrup761x25519-sha512` post-quantum key agreement algorithm is available:

Internal Only - General

```
# "$(readlink -e /usr/sbin/sshd || readlink -e /sbin/sshd)" -T | \
awk '($1=="kexalgorithms" && $2~/sntrup761x25519-sha512/) {print $2}'
```

Verify output includes `sntrup761x25519-sha512`

- IF - openSSH-server version is `9.9` or greater, run the following command to verify the `mlkem768x25519-sha256` post-quantum key agreement algorithm is also available:

```
# "$(readlink -e /usr/sbin/sshd || readlink -e /sbin/sshd)" -T | \
awk '($1=="kexalgorithms" && $2~/mlkem768x25519-sha256/) {print $2}'
```

Verify output includes `mlkem768x25519-sha256`

Remediation:

Edit or create a drop-in file under `/etc/ssh/sshd_config.d/` to enable the post-quantum cryptography key exchange algorithm(s) using the additive form. The drop-in file should be named so that lexical order places it before any other drop-in that sets `KexAlgorithms` (for example `10-cis-sshd.conf`), and the `KexAlgorithms` directive must appear above any `Include` or `Match` statements.

- IF - openSSH-server version is `9.9` or greater, enable both `sntrup761x25519-sha512` and `mlkem768x25519-sha256`:

```
# printf '%s\n' 'KexAlgorithms +sntrup761x25519-sha512,mlkem768x25519-sha256'
> /etc/ssh/sshd_config.d/10-cis-sshd.conf
```

- IF - openSSH-server version is less than `9.9`, enable only `sntrup761x25519-sha512`:

```
# printf '%s\n' 'KexAlgorithms +sntrup761x25519-sha512' >
/etc/ssh/sshd_config.d/10-cis-sshd.conf
```

Then reload (or restart) the sshd service to apply the change:

```
# systemctl reload-or-restart sshd.service
```

Verify the running configuration includes the required PQC key exchange algorithm(s):

```
# sshd -T | grep -i '^kexalgorithms'
```

> **Note:** If local site policy has stricter requirements, add any additional `KexAlgorithms` that are required by local site policy to the additive list in the example.

Run the following command to apply the configuration change to the running sshd daemon without interrupting active sessions:

```
# systemctl reload sshd
```

Default Value:

Version 9.0 - 9.8: kexalgorithms sntrup761x25519-sha512,sntrup761x25519-sha512@openssh.com,curve25519-sha256,curve25519-sha256@libssh.org,ecdh-sha2-nistp256,ecdh-sha2-nistp384,ecdh-sha2-nistp521

Internal Only - General

Version 9.9 and greater: kexalgorithms mlkem768x25519-sha256,sntrup761x25519-sha512,sntrup761x25519-sha512@openssh.com,curve25519-sha256,curve25519-sha256@libssh.org,ecdh-sha2-nistp256,ecdh-sha2-nistp384,ecdh-sha2-nistp521

References:

1. GRID: NIX-00000211

Additional Information:

OpenSSH 9.0: Introduced `sntrup761x25519-sha512` as the first post-quantum key agreement algorithm, which is a hybrid scheme combining Streamlined NTRU Prime with X25519.

OpenSSH 9.9: Added `mlkem768x25519-sha256` as another hybrid post-quantum key agreement algorithm.

OpenSSH 10.0: Made `mlkem768x25519-sha256` the default key agreement algorithm. This algorithm combines ML-KEM (a NIST-standardized key encapsulation mechanism, also known as Kyber) with the classical X25519 elliptic curve method.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.10 Encrypt Sensitive Data in Transit Encrypt sensitive data in transit. Example implementations can include: Transport Layer Security (TLS) and Open Secure Shell (OpenSSH).		●	●
v7	14.4 Encrypt All Sensitive Information in Transit Encrypt all sensitive information in transit.		●	●

Internal Only - General

5.1.15 Ensure sshd LoginGraceTime is configured (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.

More information about the openSSH server configuration is available in the "Configure SSH Server" section overview.

Rationale:

Setting the `LoginGraceTime` parameter to a low number will minimize the risk of successful brute force attacks to the SSH server. It will also limit the number of concurrent unauthenticated connections While the recommended setting is 60 seconds (1 Minute), set the number based on site policy.

Audit:

Run the following command and verify that output `LoginGraceTime` is between `1` and `60` seconds:

```
# sshd -T | grep logingracetime

logingracetime 60
```

Remediation:

Create or edit a `*.conf` file in the `/etc/ssh/sshd_config.d/` and set the `LoginGraceTime` parameter to `60` seconds or less above any `Include` entry as follows:

```
LoginGraceTime 60
```

Run the following command to apply the configuration change to the running sshd daemon without interrupting active sessions:

```
# systemctl reload sshd
```

Default Value:

LoginGraceTime 120

Internal Only - General

References:

1. GRID: NIX-00000201
2. NIST SP 800-53 Rev. 5: CM-6
3. SSHD(8)
4. SSHD_CONFIG(5)

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

5.1.16 Ensure sshd LogLevel is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

SSH provides several logging levels with varying amounts of verbosity. The `DEBUG` options are specifically not recommended other than strictly for debugging SSH communications. These levels provide so much data that it is difficult to identify important security information, and may violate the privacy of users.

More information about the openSSH server configuration is available in the "Configure SSH Server" section overview.

Rationale:

The `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.

The `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.

Audit:

Run the following command and verify that output matches `loglevel VERBOSE` or `loglevel INFO`:

```
# sshd -T | grep loglevel

loglevel VERBOSE
   - OR -
loglevel INFO
```

- IF - `Match` set statements are used in your environment, specify the connection parameters to use for the `-T` extended test mode and run the audit to verify the setting is not incorrectly configured in a match block

Example additional audit needed for a match block for the user `sshuser`:

```
# sshd -T -C user=sshuser | grep loglevel
```

Remediation:

Create or edit a `*.conf` file in the `/etc/ssh/sshd_config.d/` and set the `LogLevel` parameter to `VERBOSE` or `INFO` above any `Include` and `Match` entries as follows:

Internal Only - General

```
LogLevel VERBOSE
   - OR -
LogLevel INFO
```

Run the following command to apply the configuration change to the running sshd daemon without interrupting active sessions:

```
# systemctl reload sshd
```

Default Value:

LogLevel INFO

References:

1. GRID: NIX-00000202
2. NIST SP 800-53 Rev. 5: AU-3, AU-12, SI-5
3. https://www.ssh.com/ssh/sshd_config/

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 Collect Audit Logs Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.	●	●	●
v7	6.2 Activate audit logging Ensure that local logging has been enabled on all systems and networking devices.	●	●	●
v7	6.3 Enable Detailed Logging Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.		●	●

Internal Only - General

5.1.17 Ensure sshd MACs are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

This variable limits the types of MAC algorithms that SSH can use during communication.

More information about the openSSH server configuration is available in the "Configure SSH Server" section overview.

Rationale:

MD5 and 96-bit MAC algorithms are considered weak and have been shown to increase exploitability in SSH downgrade attacks. Weak algorithms continue to have a great deal of attention as a weak spot that can be exploited with expanded computing power. An attacker that breaks the algorithm could take advantage of a MiTM position to decrypt the SSH tunnel and capture credentials and information.

Impact:

Incorrectly configured or an overly restrictive list of MACs could lead to a "no matching MAC" error during an SSH connection. This means the client and server have no overlapping MAC algorithms which typically occurs due to a mismatch in supported MACs.

Audit:

Run the following command to verify none of the "weak" MACs are being used:

```
# sshd -T | grep -Pi -- 'macs\h+([^#\n\r]+,)?(hmac-md5|hmac-md5-96|hmac-
ripemd160|hmac-sha1-96|umac-64@openssh\.com|hmac-md5-etm@openssh\.com|hmac-
md5-96-etm@openssh\.com|hmac-ripemd160-etm@openssh\.com|hmac-sha1-96-
etm@openssh\.com|umac-64-etm@openssh\.com|umac-128-etm@openssh\.com)\b'
```

Nothing should be returned

Note: Review CVE-2023-48795 and verify the system has been patched. If the system has not been patched, review the use of the Encrypt Then Mac (etm) MACs.

The following are considered "weak" MACs, and should not be used:

Internal Only - General

```
hmac-md5
hmac-md5-96
hmac-ripemd160
hmac-sha1-96
umac-64@openssh.com
hmac-md5-etm@openssh.com
hmac-md5-96-etm@openssh.com
hmac-ripemd160-etm@openssh.com
hmac-sha1-96-etm@openssh.com
umac-64-etm@openssh.com
umac-128-etm@openssh.com
```

Remediation:

Create or edit a `*.conf` file in the `/etc/ssh/sshd_config.d/` and add/modify the `MACs` line to contain a comma separated list of the site unapproved (weak) MACs preceded with a `-` above any `Include` entries:

Example:

```
MACs -hmac-md5,hmac-md5-96,hmac-ripemd160,hmac-sha1-96,umac-
64@openssh.com,hmac-md5-etm@openssh.com,hmac-md5-96-etm@openssh.com,hmac-
ripemd160-etm@openssh.com,hmac-sha1-96-etm@openssh.com,umac-64-
etm@openssh.com,umac-128-etm@openssh.com
```

- IF - `CVE-2023-48795` has not been reviewed and addressed, the following `etm` MACs should be added to the exclude list: hmac-sha1-etm@openssh.com,hmac-sha2-256-etm@openssh.com,hmac-sha2-512-etm@openssh.com

Run the following command to apply the configuration change to the running sshd daemon without interrupting active sessions:

```
# systemctl reload sshd
```

Default Value:

MACs umac-64-etm@openssh.com,umac-128-etm@openssh.com,hmac-sha2-256-etm@openssh.com,hmac-sha2-512-etm@openssh.com,hmac-sha1-etm@openssh.com,umac-64@openssh.com,umac-128@openssh.com,hmac-sha2-256,hmac-sha2-512,hmac-sha1

References:

1. GRID: NIX-00000203
2. More information on SSH downgrade attacks can be found here: http://www.mitls.org/pages/attacks/SLOTH
3. SSHD_CONFIG(5)
4. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5
5. https://nvd.nist.gov/vuln/detail/CVE-2023-48795

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.10 Encrypt Sensitive Data in Transit Encrypt sensitive data in transit. Example implementations can include: Transport Layer Security (TLS) and Open Secure Shell (OpenSSH).		●	●
v7	14.4 Encrypt All Sensitive Information in Transit Encrypt all sensitive information in transit.		●	●
v7	16.5 Encrypt Transmittal of Username and Authentication Credentials Ensure that all account usernames and authentication credentials are transmitted across networks using encrypted channels.		●	●

Internal Only - General

5.1.18 Ensure sshd MaxAuthTries is configured (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.

More information about the openSSH server configuration is available in the "Configure SSH Server" section overview.

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 `MaxAuthTries` is `4` or less:

```
# sshd -T | grep maxauthtries

maxauthtries 4
```

- IF - `Match` set statements are used in your environment, specify the connection parameters to use for the `-T` extended test mode and run the audit to verify the setting is not incorrectly configured in a match block

Example additional audit needed for a match block for the user `sshuser`:

```
# sshd -T -C user=sshuser | grep maxauthtries
```

Remediation:

Create or edit a `*.conf` file in the `/etc/ssh/sshd_config.d/` and set the `MaxAuthTries` parameter to `4` or less above any `Include` and `Match` entries as follows:

```
MaxAuthTries 4
```

Run the following command to apply the configuration change to the running sshd daemon without interrupting active sessions:

Internal Only - General

```
# systemctl reload sshd
```

Default Value:

MaxAuthTries 6

References:

1. GRID: NIX-00000204
2. NIST SP 800-53 Rev. 5: AU-3
3. SSHD_CONFIG(5)

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	16.13 Alert on Account Login Behavior Deviation Alert when users deviate from normal login behavior, such as time-of-day, workstation location and duration.			●

Internal Only - General

5.1.19 Ensure sshd MaxSessions is configured (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.

More information about the openSSH server configuration is available in the "Configure SSH Server" section overview.

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 `MaxSessions` is `10` or less:

```
# sshd -T | grep maxsessions

maxsessions 10
```

Run the following command and verify the output:

```
grep -Psi -- '^\h*MaxSessions\h+\"?(1[1-9]|[2-9][0-9]|[1-9][0-9][0-9]+)\b'
/etc/ssh/sshd_config /etc/ssh/sshd_config.d/*.conf
```

Nothing should be returned.

- IF - `Match` set statements are used in your environment, specify the connection parameters to use for the `-T` extended test mode and run the audit to verify the setting is not incorrectly configured in a match block

Example additional audit needed for a match block for the user `sshuser`:

```
# sshd -T -C user=sshuser | grep maxsessions
```

Remediation:

Create or edit a `*.conf` file in the `/etc/ssh/sshd_config.d/` and set the `MaxSessions` parameter to `10` or less above any `Include` and `Match` entries as follows:

```
MaxSessions 10
```

Run the following command to apply the configuration change to the running sshd daemon without interrupting active sessions:

Internal Only - General

```
# systemctl reload sshd
```

Default Value:

MaxSessions 10

References:

1. GRID: NIX-00000205
2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5
3. SSHD_CONFIG(5)

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

5.1.20 Ensure sshd 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.

More information about the openSSH server configuration is available in the "Configure SSH Server" section overview.

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 to verify `MaxStartups` is `10:30:60` or more restrictive:

```
# sshd -T | awk '$1 ~ /^\s*maxstartups/{split($2, a, ":");{if(a[1] > 10 ||
a[2] > 30 || a[3] > 60) print $0}}'
```

Nothing should be returned

Remediation:

Create or edit a `*.conf` file in the `/etc/ssh/sshd_config.d/` and set the `MaxStartups` parameter to `10:30:60` or more restrictive above any `Include` entries as follows:

```
MaxStartups 10:30:60
```

Run the following command to apply the configuration change to the running sshd daemon without interrupting active sessions:

```
# systemctl reload sshd
```

Default Value:

MaxStartups 10:30:100

References:

1. GRID: NIX-00000206
2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5
3. SSHD_CONFIG(5)

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

5.1.21 Ensure sshd 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.

More information about the openSSH server configuration is available in the "Configure SSH Server" section overview.

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 to verify `PermitEmptyPasswords` is set to `no`:

```
# sshd -T | grep permitemptypasswords

permitemptypasswords no
```

- IF - `Match` set statements are used in your environment, specify the connection parameters to use for the `-T` extended test mode and run the audit to verify the setting is not incorrectly configured in a match block

Example additional audit needed for a match block for the user `sshuser`:

```
# sshd -T -C user=sshuser | grep permitemptypasswords
```

Remediation:

Create or edit a `*.conf` file in the `/etc/ssh/sshd_config.d/` and set the `PermitEmptyPasswords` parameter to `no` above any `Include` and `Match` entries as follows:

```
PermitEmptyPasswords no
```

Run the following command to apply the configuration change to the running sshd daemon without interrupting active sessions:

Internal Only - General

```
# systemctl reload sshd
```

Default Value:

PermitEmptyPasswords no

References:

1. GRID: NIX-00000207
2. NIST SP 800-53 Rev. 5: CM-1,CM-2, CM-6, CM-7, IA-5
3. SSHD_CONFIG(5)

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.2 Use Unique Passwords Use unique passwords for all enterprise assets. Best practice implementation includes, at a minimum, an 8-character password for accounts using MFA and a 14-character password for accounts not using MFA.	●	●	●
v7	4.4 Use Unique Passwords Where multi-factor authentication is not supported (such as local administrator, root, or service accounts), accounts will use passwords that are unique to that system.		●	●

Internal Only - General

5.1.22 Ensure sshd PermitRootLogin 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 `prohibit-password`.

More information about the openSSH server configuration is available in the "Configure SSH Server" section overview.

Rationale:

Disallowing `root` logins over SSH requires system admins to authenticate using their own individual account, then escalating to `root`. This limits opportunity for non-repudiation and provides a clear audit trail in the event of a security incident.

Audit:

Run the following command to verify `PermitRootLogin` is set to `no`:

```
# sshd -T | grep permitrootlogin

permitrootlogin no
```

- IF - `Match` set statements are used in your environment, specify the connection parameters to use for the `-T` extended test mode and run the audit to verify the setting is not incorrectly configured in a match block

Example additional audit needed for a match block for the user `sshuser`:

```
# sshd -T -C user=sshuser | grep permitrootlogin
```

Remediation:

Create or edit a `*.conf` file in the `/etc/ssh/sshd_config.d/` and set the `PermitRootLogin` parameter to `no` above any `Include` and `Match` entries as follows:

```
PermitRootLogin no
```

Run the following command to apply the configuration change to the running sshd daemon without interrupting active sessions:

```
# systemctl reload sshd
```

Default Value:

PermitRootLogin without-password

Internal Only - General

References:

1. GRID: NIX-00000208
2. NIST SP 800-53 Rev. 5:AC-6
3. SSHD_CONFIG(5)

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.4 Restrict Administrator Privileges to Dedicated Administrator Accounts Restrict administrator privileges to dedicated administrator accounts on enterprise assets. Conduct general computing activities, such as internet browsing, email, and productivity suite use, from the user's primary, non-privileged account.	●	●	●
v7	4.3 Ensure the Use of Dedicated Administrative Accounts 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.	●	●	●

Internal Only - General

5.1.23 Ensure sshd 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.

More information about the openSSH server configuration is available in the "Configure SSH Server" section overview.

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 to verify `PermitUserEnvironment` is set to `no`:

```
# sshd -T | grep permituserenvironment

permituserenvironment no
```

Remediation:

Create or edit a `*.conf` file in the `/etc/ssh/sshd_config.d/` and set the `PermitUserEnvironment` parameter to `no` above any `Include` entries as follows:

```
PermitUserEnvironment no
```

Run the following command to apply the configuration change to the running sshd daemon without interrupting active sessions:

```
# systemctl reload sshd
```

Default Value:

PermitUserEnvironment no

References:

1. GRID: NIX-00000209
2. NIST SP 800-53 Rev. 5: CM-1,CM-2, CM-6, CM-7, IA-5
3. SSHD(8)

Internal Only - General

4. SSHD_CONFIG(5)

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

5.1.24 Ensure sshd UsePAM is enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `UsePAM` directive enables the Pluggable Authentication Module (PAM) interface. If set to `yes` this will enable PAM authentication using `ChallengeResponseAuthentication` and `PasswordAuthentication` directives in addition to PAM account and session module processing for all authentication types.

More information about the openSSH server configuration is available in the "Configure SSH Server" section overview.

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

Audit:

Run the following command to verify `UsePAM` is set to `yes`:

```
# sshd -T | grep usepam

usepam yes
```

Remediation:

Create or edit a `*.conf` file in the `/etc/ssh/sshd_config.d/` and set the `UsePAM` parameter to `yes` above any `Include` entries as follows:

```
UsePAM yes
```

Run the following command to apply the configuration change to the running sshd daemon without interrupting active sessions:

```
# systemctl reload sshd
```

Default Value:

UsePAM yes

References:

1. GRID: NIX-00000210

Internal Only - General

2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5
3. SSHD(8)
4. SSHD_CONFIG(5)

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.2 Use Unique Passwords Use unique passwords for all enterprise assets. Best practice implementation includes, at a minimum, an 8-character password for accounts using MFA and a 14-character password for accounts not using MFA.	●	●	●
v7	4.4 Use Unique Passwords Where multi-factor authentication is not supported (such as local administrator, root, or service accounts), accounts will use passwords that are unique to that system.		●	●

Internal Only - General

5.2 Configure privilege escalation

There are various tools which allow a permitted user to execute a command as the superuser or another user, as specified by the security policy.

sudo / sudo-rs

On Debian 13 ("trixie") and Ubuntu 26.04+, the default sudo implementation is `sudo-rs` — a safety-oriented, memory-safe re-implementation of the original sudo, written in Rust. The classic `sudo` (Todd Miller's C implementation) remains available as an alternative package. Both implementations provide `/usr/bin/sudo` and read configuration from `/etc/sudoers` and files under `/etc/sudoers.d/`.

sudo documentation | sudo-rs documentation

The invoking user's real (not effective) user ID is used to determine the user name with which to query the security policy.

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

Behavioral differences between sudo-rs and classic sudo relevant to this section:

- `sudo-rs` does **not** provide a plugin architecture.
- `sudo-rs` does **not** support the `Defaults logfile=` (custom log file) directive — all sudo events are sent to syslog under the `authpriv` facility. See 5.2.3.
- `sudo-rs` has `use_pty` enabled by default; no explicit `Defaults use_pty` line is required for compliance. See 5.2.2.
- `sudo-rs` hard-codes `verifypw=all` and does not honor the `!authenticate` tag — re-authentication is always required. See 5.2.5.
- `sudo-rs` hard-codes `timestamp_type=tty`. `visiblepw` is always disabled, and `always_set_home`, `always_query_group_plugin`, `match_group_by_gid`, and `mail_badpass` are accepted but ignored. Older `sudo-rs` releases may also ignore `Defaults timestamp_timeout=N` — verify the effective value on the target with `sudo -V`. See 5.2.6.
- The `NOPASSWD` and `PASSWD` Cmnd_Spec tags and the `secure_path`, `env_reset`, and `pwfeedback` Defaults are supported by both implementations.

When a recommendation in this section depends on behavior that differs between implementations, the recommendation text identifies the applicable audit and remediation method per implementation.

pkexec

pkexec documentation

Internal Only - General

`pkexec` allows an authorized user to execute *PROGRAM* as another user. If *username* is not specified, then the program will be executed as the administrative super user, `root`.

Internal Only - General

5.2.1 Ensure sudo is installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`sudo` (and `sudo-rs`, the memory-safe Rust re-implementation that is the default on Debian 13 and Ubuntu 26.04+) 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.

The default security policy is `sudoers`, which is configured via the file `/etc/sudoers` and any entries in `/etc/sudoers.d`.

The security policy determines what privileges, if any, a user has to run `sudo`. The policy may require that users authenticate themselves with a password or another authentication mechanism. If authentication is required, `sudo` will exit if the user's password is not entered within a configurable time limit. This limit is policy-specific.

Note: sudo-rs implements a curated subset of the classic sudo `Defaults` flags and does not provide a plugin architecture. See https://manpages.debian.org/trixie/sudo-rs/sudo-rs.8.en.html for details.

Rationale:

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

The security policy determines what privileges, if any, a user has to run `sudo`. The policy may require that users authenticate themselves with a password or another authentication mechanism. If authentication is required, `sudo` will exit if the user's password is not entered within a configurable time limit. This limit is policy-specific.

Audit:

Run the following command to verify that one of `sudo-rs` or `sudo` is installed:

```
# dpkg-query -W -f='${db:Status-Status} ${Package}\n' sudo-rs sudo
2>/dev/null | awk '$1=="installed"{print $2}'
```

At least one package should be returned.

- OR -

Internal Only - General

Run the following to verify that `libsss-sudo` and `sssd` are installed (only when LDAP-backed sudoers via SSSD is required):

```
# dpkg-query -s libsss-sudo &>/dev/null && echo "libsss-sudo is installed"
# dpkg-query -s sssd &>/dev/null && echo "sssd is installed"
```

Remediation:

Install the appropriate sudo implementation for your environment.

Preferred on Debian 13 / Ubuntu 26.04+ (sudo-rs is the default):

```
# apt install sudo-rs
```

Classic sudo (if explicitly required for plugin support or non-standard Defaults flags):

```
# apt install sudo
```

Only if LDAP-backed sudoers via SSSD is required:

```
# apt install libsss-sudo sssd
```

Note: sudo-ldap is deprecated and being phased out. Debian 13 ("trixie") is the last Debian release to ship it; on Ubuntu 26.04 and later it is removed entirely and is not installable. For LDAP-backed sudoers on either distribution, use `libsss-sudo` together with `sssd`. The Debian sudo team recommends migrating existing sudo-ldap installations to `libsss-sudo` and `sssd`. See `/usr/share/doc/sudo-ldap/NEWS.Debian.gz` and Debian BTS #1033728 for more detail.

References:

1. GRID: NIX-00000215
2. SUDO(8)
3. NIST SP 800-53 Rev. 5: AC-6

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.4 Restrict Administrator Privileges to Dedicated Administrator Accounts Restrict administrator privileges to dedicated administrator accounts on enterprise assets. Conduct general computing activities, such as internet browsing, email, and productivity suite use, from the user's primary, non-privileged account.	●	●	●
v7	4.3 Ensure the Use of Dedicated Administrative Accounts 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.	●	●	●

Internal Only - General

Internal Only - General

5.2.2 Ensure sudo commands use pty (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`sudo` can be configured to run only from a pseudo terminal (`pseudo-pty`). This prevents commands launched via `sudo` from forking background processes that persist after the user's session exits.

Note: On sudo-rs (the default on Debian 13 and Ubuntu 26.04+), `use_pty` is enabled by default and an explicit `Defaults use_pty` line is not required. This recommendation is automatically satisfied on a default sudo-rs install. The audit therefore checks for two conditions: that `!use_pty` (which would disable pty allocation) is not configured anywhere, AND that the running implementation either has `use_pty` enabled by default (sudo-rs) or has `Defaults use_pty` explicitly configured (classic sudo).

Rationale:

Attackers can run a malicious program using `sudo` which would fork a background process that remains even when the main program has finished executing.

Impact:

WARNING: Editing the `sudo` configuration incorrectly can cause `sudo` to stop functioning. Always use `visudo` to modify `sudo` configuration files.

Audit:

Verify that `sudo` can only run other commands from a pseudo terminal.

1. Run the following command to verify that `!use_pty` is not set anywhere in sudoers:

```
# grep -rPi -- '^\h*Defaults\h+([^#\n\r]+,\h*)?!use_pty\b' /etc/sudoers
/etc/sudoers.d/ 2>/dev/null
```

Nothing should be returned.

2. Verify `use_pty` is explicitly configured in sudoers:

```
# grep -rPi -- '^\h*Defaults\h+([^#\n\r]+,\h*)?use_pty\b' /etc/sudoers
/etc/sudoers.d/ 2>/dev/null
```

At least one line such as `Defaults use_pty` should be returned.

Internal Only - General

Note: On sudo-rs (the default on Debian 13 and Ubuntu 26.04+), `use_pty` is enabled by the implementation default. An explicit `Defaults use_pty` entry is still required for benchmark compliance and portability — this ensures the setting is documented and resilient to future changes in default behavior.

Remediation:

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

```
Defaults use_pty
```

Edit the file `/etc/sudoers` with `visudo` and any files in `/etc/sudoers.d/` with `visudo -f <PATH TO FILE>` and remove any occurrence of `!use_pty`

Note:

- sudo will read each file in `/etc/sudoers.d`, skipping file names that end in `~` or contain a `.` character to avoid causing problems with package manager or editor temporary/backup files.
- Files are parsed in sorted lexical order. That is, `/etc/sudoers.d/01_first` will be parsed before `/etc/sudoers.d/10_second`.
- Be aware that because the sorting is lexical, not numeric, `/etc/sudoers.d/1_whoops` would be loaded after `/etc/sudoers.d/10_second`.
- Using a consistent number of leading zeroes in the file names can be used to avoid such problems.

References:

1. GRID: NIX-00000214
2. SUDO(8)
3. sudoers(5)
4. VISUDO(8)
5. NIST SP 800-53 Rev. 5: AC-6

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.4 Restrict Administrator Privileges to Dedicated Administrator Accounts 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.	●	●	●

Internal Only - General

Controls Version	Control	IG 1	IG 2	IG 3
v7	4.3 Ensure the Use of Dedicated Administrative Accounts 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.	●	●	●

Internal Only - General

5.2.3 Ensure sudo log file exists (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`sudo` logs successful and unsuccessful privilege-escalation attempts. On systems running sudo-rs (the default on Debian 13 and Ubuntu 26.04+), the `logfile` (`log_file`) Defaults option is **not supported** — sudo-rs always sends events to syslog under the `authpriv` facility. On classic sudo, both a dedicated `Defaults logfile=` file and syslog capture are available.

This recommendation ensures sudo events are captured in a durable, auditable log location regardless of the implementation in use.

Rationale:

A durable record of sudo invocations simplifies forensic investigation of privilege-escalation events, supports compliance auditing, and provides a tamper-evidence trail when paired with centralized log collection. On sudo-rs this is achieved exclusively via syslog/journald; on classic sudo, either a Defaults logfile or syslog can be used.

Impact:

WARNING: Editing the `sudo` configuration incorrectly can cause `sudo` to stop functioning. Always use `visudo` to modify `sudo` configuration files.

Audit:

Determine which sudo implementation is in use:

```
# dpkg-query -s sudo-rs 2>/dev/null | grep -q '^Status: install ok installed'
&& echo "sudo-rs" || echo "classic sudo"
```

Example Output

```
sudo-rs
```

Method A — sudo-rs (or any system where syslog capture is desired): syslog/journald logging check Verify sudo events are reaching syslog/journald. Generate a recent event and check:

```
# journalctl -t sudo -t sudo-rs --since=-5min --no-pager | head -20
```

Verify the output contains at least one sudo event for the test invocation.

rsyslog logging check Verify sudo events are reaching rsyslog. Generate a recent event and check:

Internal Only - General

```
# grep -E '\b(sudo|sudo-rs)\[[0-9]+\]:' /var/log/auth.log 2>/dev/null | tail
-5
```

Verify the output contains at least one sudo event for the test invocation.

Method B — classic sudo only (logfile= Defaults are supported):

```
# grep -rPsi
"^\h*Defaults\h+([^#]+,\h*)?logfile\h*=\h*(\"|')?\H+(\"|')?(,\h*\H+\h*)*\h*(#
.*)?$" /etc/sudoers /etc/sudoers.d/ 2>/dev/null
```

Verify the output matches:

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

Pass condition: Method A succeeds on any system, OR Method B succeeds on a classic-sudo system.

Remediation:

Configure durable capture of sudo events using the method appropriate for the installed implementation.

Method A — sudo-rs (Debian 13 default) and classic sudo via syslog:

Ensure rsyslog or systemd-journald is capturing `authpriv.*` events. To route sudo events to a dedicated file via rsyslog:

```
# printf '%s\n' '# CIS 5.2.3 — durable sudo event capture' \
  'authpriv.*    /var/log/sudo.log' \
  > /etc/rsyslog.d/30-sudo.conf
# systemctl restart rsyslog
```

Then ensure the resulting file has appropriate permissions:

```
# touch /var/log/sudo.log
# chmod 0640 /var/log/sudo.log
# chown root:adm /var/log/sudo.log
```

Method B — classic sudo only (logfile= Defaults are supported and may be preferred):

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

WARNING: `Defaults logfile="..."` is silently ignored by sudo-rs. Do not rely on this method on a sudo-rs system — use Method A instead.

Notes:

- sudo will read each file in `/etc/sudoers.d`, skipping file names that end in `~` or contain a `.` character to avoid causing problems with package manager or editor temporary/backup files.

Internal Only - General

- Files are parsed in sorted lexical order. That is, `/etc/sudoers.d/01_first` will be parsed before `/etc/sudoers.d/10_second`.
- Be aware that because the sorting is lexical, not numeric, `/etc/sudoers.d/1_whoops` would be loaded after `/etc/sudoers.d/10_second`.
- Using a consistent number of leading zeroes in the file names can be used to avoid such problems.

References:

1. GRID: NIX-00000216
2. SUDO(8)
3. sudoers(5)
4. VISUDO(8)
5. NIST SP 800-53 Rev. 5: AU-3, AU-12

Additional Information:

visudo edits the sudoers file in a safe fashion, analogous to vipw(8). visudo locks the sudoers file against multiple simultaneous edits, provides basic sanity checks, and checks for parse errors. If the sudoers file is currently being edited you will receive a message to try again later.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 Collect Detailed Audit Logs Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.		●	●
v7	6.3 Enable Detailed Logging Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.		●	●

Internal Only - General

5.2.4 Ensure users must provide password for escalation (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

The operating system must be configured so that users must provide a password for privilege escalation.

Rationale:

Without (re-)authentication, users may access resources or perform tasks for which they do not have authorization.

When operating systems provide the capability to escalate a functional capability, it is critical the user (re-)authenticate.

Impact:

This will prevent automated processes from being able to elevate privileges.

Audit:

Note: If passwords are not being used for authentication, this is not applicable.

Verify the operating system requires users to supply a password for privilege escalation.

Check the configuration of `/etc/sudoers` and the files in `/etc/sudoers.d/` for any uncommented use of the `NOPASSWD` tag:

```
# grep -rPi -- '^\s*[^#].*\bNOPASSWD\s*:' /etc/sudoers /etc/sudoers.d/
2>/dev/null
```

If any line is returned, refer to the remediation procedure below.

Note: The `NOPASSWD` tag is supported on both classic sudo and sudo-rs; the audit logic is identical for both implementations.

Remediation:

Based on the outcome of the audit procedure, use `visudo -f <PATH TO FILE>` to edit the relevant sudoers file.

Remove any line with occurrences of `NOPASSWD` tags in the file.

Internal Only - General

References:

1. GRID: NIX-00000218
2. NIST SP 800-53 Rev. 5: AC-6

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.4 Restrict Administrator Privileges to Dedicated Administrator Accounts Restrict administrator privileges to dedicated administrator accounts on enterprise assets. Conduct general computing activities, such as internet browsing, email, and productivity suite use, from the user's primary, non-privileged account.	●	●	●
v7	4.3 Ensure the Use of Dedicated Administrative Accounts 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.	●	●	●

Internal Only - General

5.2.5 Ensure re-authentication for privilege escalation is not disabled globally (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The operating system must be configured so that users must re-authenticate for privilege escalation.

Note: On sudo-rs (the default on Debian 13 and Ubuntu 26.04+), `verifypw` is hard-coded to `all` and the `!authenticate` tag does not bypass authentication. This recommendation is automatically enforced on sudo-rs even if a `!authenticate` tag is present in sudoers files. Explicit removal of `!authenticate` remains the policy goal for portability and clarity (the same configuration will fail on a classic-sudo system).

Rationale:

Without re-authentication, users may access resources or perform tasks for which they do not have authorization.

When operating systems provide the capability to escalate a functional capability, it is critical the user re-authenticate.

Audit:

Verify the operating system requires users to re-authenticate for privilege escalation.

Check the configuration of `/etc/sudoers` and the files in `/etc/sudoers.d/` for any uncommented use of the `!authenticate` tag:

```
# grep -rPi -- '^\s*[^#].*!authenticate\b' /etc/sudoers /etc/sudoers.d/
2>/dev/null
```

If any line is returned, refer to the remediation procedure below.

Remediation:

Configure the operating system to require users to reauthenticate for privilege escalation.

Based on the outcome of the audit procedure, use `visudo -f <PATH TO FILE>` to edit the relevant sudoers file.

Remove any occurrences of `!authenticate` tags in the file(s).

Internal Only - General

References:

1. GRID: NIX-00000213
2. NIST SP 800-53 Rev. 5: AC-6

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.4 Restrict Administrator Privileges to Dedicated Administrator Accounts Restrict administrator privileges to dedicated administrator accounts on enterprise assets. Conduct general computing activities, such as internet browsing, email, and productivity suite use, from the user's primary, non-privileged account.	●	●	●
v7	4.3 Ensure the Use of Dedicated Administrative Accounts 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.	●	●	●

Internal Only - General

5.2.6 Ensure sudo timestamp_timeout is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

sudo `timestamp_timeout` controls how long a user's sudo privileges remain active after the initial password entry.

Rationale:

A timeout value reduces the window of opportunity for unauthorized privileged sudo access.

Audit:

Run the following command to verify that `timestamp_timeout` is explicitly configured:

```
# grep -rPsi -- '^\h*Defaults\h+([^#\n\r]+,\h*)?timestamp_timeout\h*=\h*[-0-9]+' \
    /etc/sudoers /etc/sudoers.d/ 2>/dev/null
```

Example Output:

```
/etc/sudoers:Defaults env_reset, timestamp_timeout=15
```

Verify the output meets all of the following:

- **Output is returned.** If nothing is returned, `timestamp_timeout` is not explicitly configured — refer to the remediation procedure.
- **No returned value is negative.** A negative value (e.g. `timestamp_timeout=-1`) disables the authentication timeout entirely.
- **No returned value is greater than `15`.** Values exceeding 15 minutes do not meet the benchmark requirement.

Remediation:

Edit `/etc/sudoers` or a file in `/etc/sudoers.d/` with `visudo -f <PATH TO FILE>` and set `timestamp_timeout=` to 15 minutes or less per your site policy. The value is in minutes. This entry may appear on its own or on the same line as other Defaults such as `env_reset`.

Example 1:

```
Defaults env_reset, timestamp_timeout=15
```

Example 2:

Internal Only - General

```
Defaults timestamp_timeout=15
Defaults env_reset
```

Note: On older sudo-rs releases the `timestamp_timeout` directive may be ignored. After applying the remediation on a sudo-rs target, run `sudo -V` and confirm the reported timeout matches expectations. If the value is not honored, upgrade sudo-rs to a release that supports `timestamp_timeout` or apply a compensating control via site policy.

Default Value:

Not explicitly configured. Classic sudo compiled-in default: 5 minutes. sudo-rs default: implementation-defined (varies by release). Explicit configuration is required by this benchmark.

References:

1. GRID: NIX-00000217
2. https://www.sudo.ws/man/1.9.0/sudoers.man.html
3. NIST SP 800-53 Rev. 5: AC-6

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.4 Restrict Administrator Privileges to Dedicated Administrator Accounts Restrict administrator privileges to dedicated administrator accounts on enterprise assets. Conduct general computing activities, such as internet browsing, email, and productivity suite use, from the user's primary, non-privileged account.	●	●	●
v7	4.3 Ensure the Use of Dedicated Administrative Accounts 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.	●	●	●

Internal Only - General

5.2.7 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 is largely superseded by `sudo` (and on Debian 13 / Ubuntu 26.04+, the `sudo-rs` Rust re-implementation), which allows for more granular control over privileged access. Normally, the `su` command can be executed by any user. By uncommenting the `pam_wheel.so` statement in `/etc/pam.d/su`, the `su` command will only allow users in a specific group to execute `su`. This group should be empty to reinforce the use of `sudo` for privileged access.

Rationale:

Restricting the use of `su` , and using `sudo` in its place, provides system administrators better control of the escalation of user privileges to execute privileged commands. The sudo utility also provides a better logging and audit mechanism, as it can log each command executed via `sudo` , whereas `su` can only record that a user executed the `su` program.

Audit:

Run the following command to verify that `pam_wheel.so` is configured with both the `use_uid` option and a `group=` argument:

```
# grep -Pi
'^\s*auth\s+(?:required|requisite)\s+pam_wheel\.so(?=.*\buse_uid\b)(?=.*\bgro
up=\S+).*' /etc/pam.d/su
```

Verify the output matches (the order of `use_uid` and `group=` does not matter):

```
auth required pam_wheel.so use_uid group=<group_name>
```

Then run the following command and verify that the group named in `group=<group_name>` contains no users:

```
# grep "^<group_name>:" /etc/group
```

Verify the trailing user list field is empty (a non-empty list of members indicates the rule will allow those members to use `su`):

```
<group_name>:x:<GID>:
```

Remediation:

Create an empty group that will be specified for use of the `su` command. The group should be named according to site policy.

Internal Only - General

Example:

```
# groupadd sugroup
```

Add the following line to the `/etc/pam.d/su` file, specifying the empty group:

```
auth required pam_wheel.so use_uid group=sugroup
```

References:

1. GRID: NIX-00000212
2. NIST SP 800-53 Rev. 5: AC-3. MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 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.	●	●	●

5.3 Pluggable Authentication Modules

Pluggable Authentication Modules (PAM) is a service that implements modular authentication modules on Linux 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.

Note: This section includes guidance that requires the additional modules and the latest version of software packages as covered in the "Configure required security packages" subsection of the "Package Management" section

Internal Only - General

5.3.1 Configure PAM software packages

Updated versions of PAM include additional functionality

Internal Only - General

5.3.1.1 Ensure latest version of pam is installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Linux Pluggable Authentication Modules (PAM) is a suite of libraries that allow a Linux system administrator to configure methods to authenticate users. It provides a flexible and centralized way to switch authentication methods for secured applications by using configuration files instead of changing application code

The `libpam-runtime` provides the runtime support for the PAM library

Rationale:

Older versions of the `libpam-runtime` package may not include the latest security and feature patches and updates.

Note: This Benchmark includes Recommendations that depend on newer `libpam-runtime` features.

Audit:

Run the following command to verify `libpam-runtime` is installed:

```
# dpkg-query -s libpam-runtime &>/dev/null && echo "libpam-runtime is
installed"

libpam-runtime is installed
```

Run the following command to verify `libpam-runtime` is the latest version:

```
# apt list --upgradable 2>&1 | grep -P '^libpam-runtime\b'
```

Nothing should be returned

Remediation:

Run the following command to install the latest version of `libpam-runtime`:

```
# apt install libpam-runtime
```

References:

1. GRID: NIX-00000222
2. PAM(7)

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

5.3.1.2 Ensure latest version of libpam-modules is installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`libpam-modules` is a package containing a set of Pluggable Authentication Modules (PAM) which allows system administrators to configure different authentication methods for user logins, providing flexibility in how users can access applications by using various authentication modules to include password checks and other security mechanisms.

Rationale:

Older versions of the `libpam-modules` package may not include the latest security and feature patches and updates.

Note: This Benchmark includes Recommendations that depend on newer `libpam-modules` features.

Audit:

Run the following command to verify `libpam-modules` is installed:

```
# dpkg-query -s libpam-modules &>/dev/null && echo "libpam-modules is
installed"

libpam-modules is installed
```

Run the following command to verify `libpam-modules` is the latest version:

```
# apt list --upgradable 2>&1 | grep -P '^libpam-modules\b'
```

Nothing should be returned

Remediation:

Run the following command to install the latest version of `libpam-modules`:

```
# apt install libpam-modules
```

References:

1. GRID: NIX-00000220

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

5.3.1.3 Ensure latest version of libpam-pwquality is installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`libpwquality` provides common functions for password quality checking and scoring them based on their apparent randomness. The library also provides a function for generating random passwords with good pronounceability.

This module can be plugged into the password stack of a given service to provide some plug-in strength-checking for passwords. The code was originally based on `pam_cracklib` module and the module is backwards compatible with its options.

Rationale:

Strong passwords reduce the risk of systems being hacked through brute force methods.

Older versions of the `libpam-pwquality` package may not include the latest security and feature patches and updates.

Audit:

Run the following command to verify `libpam-pwquality` is installed:

```
# dpkg-query -s libpam-pwquality &>/dev/null && echo "libpam-pwquality is
installed"

libpam-pwquality is installed
```

Run the following command to verify `libpam-pwquality` is the latest version:

```
# apt list --upgradable 2>&1 | grep -P '^libpam-pwquality\b'
```

Nothing should be returned

Remediation:

Run the following command to install the latest version of `libpam-pwquality`:

```
# apt install libpam-pwquality
```

References:

1. GRID: NIX-00000221
2. https://packages.debian.org/buster/libpam-pwquality

Internal Only - General

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.		●	●

Internal Only - General

5.3.1.4 Ensure latest version of cracklib-runtime is installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`cracklib-runtime` provides runtime support for password checker library cracklib2.

Rationale:

Strong passwords reduce the risk of systems being hacked through brute force methods.

Older versions of the `cracklib-runtime` package may not include the latest security and feature patches and updates.

Audit:

Run the following command to verify `cracklib-runtime` is installed:

```
# dpkg-query -s cracklib-runtime&>/dev/null && echo "cracklib-runtime is
installed"

cracklib-runtime is installed
```

Run the following command to verify `cracklib-runtime` is the latest version:

```
# apt list --upgradable 2>&1 | grep -P '^cracklib-runtime\b'
```

Nothing should be returned

Remediation:

Run the following command to install the latest version of `cracklib-runtime`:

```
# apt install cracklib-runtime
```

References:

1. GRID: NIX-00000395
2. https://packages.ubuntu.com/questing/cracklib-runtime

Internal Only - General

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.		●	●

Internal Only - General

5.3.2 Configure pam-auth-update profiles

`pam-auth-update` is a utility that permits configuring the central authentication policy for the system using pre-defined profiles as supplied by PAM module packages.

Profiles - Shipped in the `/usr/share/pam-configs/` directory specify the modules, with options, to enable; the preferred ordering with respect to other profiles; and whether a profile should be enabled by default. Packages providing PAM modules register their profiles at install time by calling pam-auth-update --package.

Selection of profiles is done using the standard `debconf` interface. The profile selection question will be asked at `medium` priority when packages are added or removed, so no user interaction is required by default. Users may invoke `pam-auth-update` directly to change their authentication configuration.

The `pam-auth-update` script makes every effort to respect local changes to `/etc/pam.d/common-*.` Local modifications to the list of module options will be preserved, and additions of modules within the managed portion of the stack will cause pam-auth-update to treat the config files as locally modified and not make further changes to the config files unless given the `--force` option.

If the user specifies that `pam-auth-update` should override local configuration changes, the locally-modified files will be saved in /etc/pam.d/ with a suffix of `.pam-old`.

Internal Only - General

5.3.2.1 Ensure pam_unix module is enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`pam_unix` is the standard Unix authentication module. It uses standard calls from the system's libraries to retrieve and set account information as well as authentication. Usually this is obtained from the `/etc/passwd` and if shadow is enabled, the `/etc/shadow` file as well.

The account component performs the task of establishing the status of the user's account and password based on the following shadow elements: `expire`, `last_change`, `max_change`, `min_change`, `warn_change`. In the case of the latter, it may offer advice to the user on changing their password or, through the `PAM_AUTHTOKEN_REQD` return, delay giving service to the user until they have established a new password. The entries listed above are documented in the shadow(5) manual page. Should the user's record not contain one or more of these entries, the corresponding shadow check is not performed.

The authentication component performs the task of checking the users credentials (password). The default action of this module is to not permit the user access to a service if their official password is blank.

Rationale:

The system should only provide access after performing authentication of a user.

Audit:

Run the following command to verify that `pam_unix` is enabled:

```
# grep -PH -- '\bpam_unix\.so\b' /etc/pam.d/common-
{account,auth,password,session,session-noninteractive}
```

Output should be similar to:

Internal Only - General

```
/etc/pam.d/common-account:account [success=1 new_authtok_reqd=done
default=ignore] pam_unix.so
/etc/pam.d/common-auth:auth [success=2 default=ignore] pam_unix.so
try_first_pass
/etc/pam.d/common-password:password [success=1 default=ignore] pam_unix.so
obscure use_authtok try_first_pass yescrypt
/etc/pam.d/common-session:session required pam_unix.so
/etc/pam.d/common-session-noninteractive:session required pam_unix.so
```

Remediation:

Run the following command to enable the `pam_unix` module:

```
# pam-auth-update --enable unix
```

Note: If a site specific custom profile is being used in your environment to configure PAM that includes the configuration for the `pam_unix` module, enable that module instead

References:

1. GRID: NIX-00000231
2. NIST SP 800-53 Rev. 5: IA-5(1)

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.	●	●	●

Internal Only - General

5.3.2.2 Ensure pam_faillock module is enabled (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 the configured number of consecutive failed authentications (this is defined by the `deny` parameter in the faillock configuration). It stores the failure records into per-user files in the tally directory.

Rationale:

Locking out user IDs after n unsuccessful consecutive login attempts mitigates brute force password attacks against your systems.

Impact:

If `pam_faillock` is not enabled, failed login attempts will not be tracked and accounts will not be locked after repeated failures, leaving the system vulnerable to brute-force attacks.

Audit:

Run the following commands to verify that `pam_faillock` is enabled:

```
# grep -P -- '\bpam_faillock\.so\b' /etc/pam.d/common-{auth,account}
```

Output should be similar to:

```
/etc/pam.d/common-auth:auth     requisite
pam_faillock.so preauth
/etc/pam.d/common-auth:auth     [default=die]
pam_faillock.so authfail
/etc/pam.d/common-account:account       required
pam_faillock.so
```

Remediation:

Create two pam-auth-update profiles in `/usr/share/pam-configs/`:

1. Create the `faillock` profile in `/usr/share/pam-configs/` with the following lines:

Internal Only - General

```
Name: Enable pam_faillock to deny access
Default: yes
Priority: 0
Auth-Type: Primary
Auth:
        [default=die]                   pam_faillock.so authfail
```

Example Script:

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

{
   arr=('Name: Enable pam_faillock to deny access' 'Default: yes' 'Priority:
0' 'Auth-Type: Primary' 'Auth:' '        [default=die]
pam_faillock.so authfail')
   printf '%s\n' "${arr[@]}" > /usr/share/pam-configs/faillock
}
```

2. Create the `faillock_notify` profile in `/usr/share/pam-configs/` with the following lines:

```
Name: Notify of failed login attempts and reset count upon success
Default: yes
Priority: 1024
Auth-Type: Primary
Auth:
        requisite                       pam_faillock.so preauth
Account-Type: Primary
Account:
        required                        pam_faillock.so
```

Example Script:

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

{
   arr=('Name: Notify of failed login attempts and reset count upon success'
'Default: yes' 'Priority: 1024' 'Auth-Type: Primary' 'Auth:' '
requisite                       pam_faillock.so preauth' 'Account-Type:
Primary' 'Account:' '        required
pam_faillock.so')
   printf '%s\n' "${arr[@]}" > /usr/share/pam-configs/faillock_notify
}
```

Run the following command to update the `common-auth` and `common-account` PAM files with the new profiles:

```
# pam-auth-update --enable <profile_filename>
```

Example:

```
# pam-auth-update --enable faillock
# pam-auth-update --enable faillock_notify
```

Note:

Internal Only - General

- The name used for the file must be used in the `pam-auth-update --enable` command
- The `Name:` line should be easily recognizable and understood
- The `Priority:` Line is important as it effects the order of the lines in the `/etc/pam.d/` files
- If a site specific custom profile is being used in your environment to configure PAM that includes the configuration for the `pam_faillock` module, enable that module instead

References:

1. GRID: NIX-00000223

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	6.2 Establish an Access Revoking Process Establish and follow a process, preferably automated, for revoking access to enterprise assets, through disabling accounts immediately upon termination, rights revocation, or role change of a user. Disabling accounts, instead of deleting accounts, may be necessary to preserve audit trails.	●	●	●
v7	16.7 Establish Process for Revoking Access Establish and follow an automated process for revoking system access by disabling accounts immediately upon termination or change of responsibilities of an employee or contractor . Disabling these accounts, instead of deleting accounts, allows preservation of audit trails.		●	●

Internal Only - General

5.3.2.3 Ensure pam_pwquality module is enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `pam_pwquality.so` module performs password quality checking. This module can be plugged into the password stack of a given service to provide strength-checking for passwords. The code was originally based on pam_cracklib module and the module is backwards compatible with its options.

The action of this module is to prompt the user for a password and check its strength against a system dictionary and a set of rules for identifying poor choices.

The first action is to prompt for a single password, check its strength and then, if it is considered strong, prompt for the password a second time (to verify that it was typed correctly on the first occasion). All being well, the password is passed on to subsequent modules to be installed as the new authentication token.

Rationale:

Use of a unique, complex passwords helps to increase the time and resources required to compromise the password.

Impact:

Without `pam_pwquality` enabled, password strength checks will not be enforced. Users will be able to set weak passwords that are easily guessed or cracked.

Audit:

Run the following command to verify that `pam_pwquality.so` is enabled:

```
# grep -P -- '\bpam_pwquality\.so\b' /etc/pam.d/common-password
```

Output should be similar to:

```
password   requisite   pam_pwquality.so retry=3
```

Remediation:

Run the following script to verify the `pam_pwquality.so` line exists in a `pam-auth-update` profile:

```
# grep -P -- '\bpam_pwquality\.so\b' /usr/share/pam-configs/*
```

Output should be similar to:

Internal Only - General

```
/usr/share/pam-configs/pwquality: requisite pam_pwquality.so retry=3
/usr/share/pam-configs/pwquality: requisite pam_pwquality.so retry=3
```

- IF - similar output is returned:

Run the following command to update `/etc/pam.d/common-password` with the returned profile:

```
# pam-auth-update --enable {PROFILE_NAME}
```

Example:

```
# pam-auth-update --enable pwquality
```

- IF - similar output is **NOT** returned:

Create a pam-auth-update profile in `/usr/share/pam-configs/` with the following lines:

```
Name: Pwquality password strength checking
Default: yes
Priority: 1024
Conflicts: cracklib
Password-Type: Primary
Password:
 requisite pam_pwquality.so retry=3
```

Example:

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

{
 arr=('Name: Pwquality password strength checking' 'Default: yes' 'Priority:
1024' 'Conflicts: cracklib' 'Password-Type: Primary' 'Password:' ' requisite
pam_pwquality.so retry=3')
 printf '%s\n' "${arr[@]}" > /usr/share/pam-configs/pwquality
}
```

Run the following command to update `/etc/pam.d/common-password` with the `pwquality` profile:

```
# pam-auth-update --enable pwquality
```

Note:

- The name used for the file must be used in the `pam-auth-update --enable` command
- The `Name:` line should be easily recognizable and understood
- The `Priority:` Line is important as it effects the order of the lines in the `/etc/pam.d/` files
- If a site specific custom profile is being used in your environment to configure PAM that includes the configuration for the `pam_pwquality` module, enable that module instead

Internal Only - General

References:

1. GRID: NIX-00000226

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.		●	●

Internal Only - General

5.3.2.4 Ensure pam_pwhistory module is enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `pam_pwhistory.so` module saves the last passwords for each user in order to force password change history and keep the user from alternating between the same password too frequently.

This module does not work together with kerberos. In general, it does not make much sense to use this module in conjunction with `NIS` or `LDAP`, since the old passwords are stored on the local machine and are not available on another machine for password history checking.

Rationale:

Use of a unique, complex passwords helps to increase the time and resources required to compromise the password.

Impact:

Without `pam_pwhistory` enabled, the system will not enforce password history, allowing users to reuse previous passwords immediately after a change.

Audit:

Run the following command to verify that `pam_pwhistory.so` is enabled:

```
# grep -P -- '\bpam_pwhistory\.so\b' /etc/pam.d/common-password
```

Output should be similar to:

```
password   requisite   pam_pwhistory.so remember=24 enforce_for_root
use_authtok
```

Remediation:

Run the following script to verify the `pam_pwhistory.so` line exists in a `pam-auth-update` profile:

```
# grep -P -- '\bpam_pwhistory\.so\b' /usr/share/pam-configs/*
```

Output should be similar to:

```
/usr/share/pam-configs/pwhistory: requisite pam_pwhistory.so remember=24
enforce_for_root use_authtok
```

- IF - similar output is returned:

Internal Only - General

Run the following command to update `/etc/pam.d/common-password` with the returned profile:

```
# pam-auth-update --enable {PROFILE_NAME}
```

Example:

```
# pam-auth-update --enable pwhistory
```

- IF - similar output is **NOT** returned:

Create a `pwhistory` profile in `/usr/share/pam-configs/` with the following lines:

```
Name: pwhistory password history checking
Default: yes
Priority: 1024
Password-Type: Primary
Password: requisite pam_pwhistory.so remember=24 enforce_for_root use_authtok
```

Example Script:

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

{
 arr=('Name: pwhistory password history checking' 'Default: yes' 'Priority:
1024' 'Password-Type: Primary' 'Password:' ' requisite pam_pwhistory.so
remember=24 enforce_for_root use_authtok')
 printf '%s\n' "${arr[@]}" > /usr/share/pam-configs/pwhistory
}
```

Run the following command to update `/etc/pam.d/common-password` with the `pwhistory` profile:

```
# pam-auth-update --enable pwhistory
```

Note:

- The name used for the file must be used in the `pam-auth-update --enable` command
- The `Name:` line should be easily recognizable and understood
- The `Priority:` Line is important as it effects the order of the lines in the `/etc/pam.d/` files
- If a site specific custom profile is being used in your environment to configure PAM that includes the configuration for the `pam_pwhistory` module, enable that module instead

If the audit identified a non-compliant or conflicting `pam_pwhistory` entry directly in `/etc/pam.d/common-password`, run the following:

1. Check for a conflicting `pam_pwhistory` entry in `common-password`:

Internal Only - General

```
# grep -P '^\s*password\s+' /etc/pam.d/common-password | grep pam_pwhistory
```

- IF - a non-compliant entry is found (e.g., missing `use_authtok`, incorrect `remember` count, or wrong control flag), remove it and re-enable via `pam-auth-update`:

```
# sed -i '/pam_pwhistory/d' /etc/pam.d/common-password
# pam-auth-update --enable pwhistory
```

- IF - no `pam_pwhistory` entry exists in `common-password`, no direct action is needed; the `pam-auth-update --enable` step above will populate the entry correctly.

References:

1. GRID: NIX-00000225

Additional Information:

Ticket #27009: Add remediation steps for handling a conflicting or non-compliant pam_pwhistory entry in /etc/pam.d/common-password, complementing the existing pam-auth-update guidance.

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.		●	●

5.3.3 Configure PAM Arguments

Pluggable Authentication Modules (PAM) uses arguments to pass information to a pluggable module during authentication for a particular module type. These arguments allow the PAM configuration files for particular programs to use a common PAM module but in different ways.

Invalid arguments are ignored and do not otherwise affect the success or failure of the PAM module. When an invalid argument is passed, an error is usually written to `/var/log/messages` file. However, since the reporting method is controlled by the PAM module, the module must be written correctly to log the error to this file.

Internal Only - General

Note: If custom PAM files are being used, for this section's remediation, the corresponding files in `/etc/pam.d/` would need to be edited directly, and the `pam-auth-update --enable <EDITED_PROFILE_NAME>` command skipped

Internal Only - General

5.3.3.1 Configure pam_faillock module

`pam_faillock.so` provides a way to configure the default settings for locking the user after multiple failed authentication attempts.

Options:

- `<dir=/path/to/tally-directory>` - The directory where the user files with the failure records are kept. The default is /var/run/faillock. Note: These files will disappear after reboot on systems configured with directory /var/run/faillock mounted on virtual memory.
- `audit` - Will log the user name into the system log if the user is not found.
- `silent` - Don't print informative messages to the user. Please note that when this option is not used there will be difference in the authentication behavior for users which exist on the system and non-existing users.
- `no_log_info` - Don't log informative messages via syslog(3).
- `local_users_only` - Only track failed user authentications attempts for local users in /etc/passwd and ignore centralized (AD, IdM, LDAP, etc.) users. The faillock(8) command will also no longer track user failed authentication attempts. Enabling this option will prevent a double-lockout scenario where a user is locked out locally and in the centralized mechanism.
- `nodelay` - Don't enforce a delay after authentication failures.
- `deny=<n>` - Deny access if the number of consecutive authentication failures for this user during the recent interval exceeds *<n>*. The default is 3.
- `fail_interval=n` - The length of the interval during which the consecutive authentication failures must happen for the user account lock out is n seconds. The default is 900 (15 minutes).
- `unlock_time=n` - The access will be re-enabled after n seconds after the lock out. The value 0 has the same meaning as value never - the access will not be re-enabled without resetting the faillock entries by the faillock(8) command. The default is 600 (10 minutes). Note that the default directory that pam_faillock uses is usually cleared on system boot so the access will be also re-enabled after system reboot. If that is undesirable a different tally directory must be set with the dir option. Also note that it is usually undesirable to permanently lock out users as they can become easily a target of denial of service attack unless the usernames are random and kept secret to potential attackers.
- `even_deny_root` - Root account can become locked as well as regular accounts.
- `root_unlock_time=n` - This option implies even_deny_root option. Allow access after n seconds to root account after the account is locked. In case the option is not specified the value is the same as of the unlock_time option.
- `admin_group=name` - If a group name is specified with this option, members of the group will be handled by this module the same as the root account (the options even_deny_root and root_unlock_time will apply to them. By default the option is not set.

Internal Only - General

5.3.3.1.1 Ensure password failed attempts lockout is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `deny=<n>` option will deny access if the number of consecutive authentication failures for this user during the recent interval exceeds <*n*>.

Rationale:

Locking out user IDs after *n* unsuccessful consecutive login attempts mitigates brute force password attacks against your systems.

Impact:

Locking out accounts after repeated failed attempts can be used to deny service to legitimate users if an attacker submits multiple incorrect authentication attempts. Consider combining with `root_unlock_time` (5.3.3.1.3) to limit the denial-of-service exposure for the root account.

Audit:

Run the following command to verify that Number of failed logon attempts before the account is locked is no greater than `5` and meets local site policy:

```
# grep -Pi -- '^\h*deny\h*=\h*[1-5]\b' /etc/security/faillock.conf

deny = 5
```

Run the following command to verify that the `deny` argument has not been set, or `5` or less and meets local site policy:

```
# grep -Pi --
'^\h*auth\h+(requisite|required|sufficient)\h+pam_faillock\.so\h+([^#\n\r]+\h
+)?deny\h*=\h*(0|[6-9]|[1-9][0-9]+)\b' /etc/pam.d/common-auth
```

Nothing should be returned

Remediation:

Create or edit the following line in `/etc/security/faillock.conf` setting the `deny` option to `5` or less:

```
deny = 5
```

Run the following command:

Internal Only - General

```
# grep -Pl -- '\bpam_faillock\.so\h+([^#\n\r]+\h+)?deny\b' /usr/share/pam-configs/*
```

Edit any returned files and remove the `deny=<N>` arguments from the `pam_faillock.so` line(s).

Default Value:

deny = 3

References:

1. GRID: NIX-00000235

Additional Information:

If a user has been locked out because they have reached the maximum consecutive failure count defined by `deny=` in the `pam_faillock.so` module, the user can be unlocked by issuing the command `faillock --user <USERNAME> --reset`. This command sets the failed count to 0, effectively unlocking the user.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	6.2 Establish an Access Revoking Process Establish and follow a process, preferably automated, for revoking access to enterprise assets, through disabling accounts immediately upon termination, rights revocation, or role change of a user. Disabling accounts, instead of deleting accounts, may be necessary to preserve audit trails.	●	●	●
v7	16.7 Establish Process for Revoking Access Establish and follow an automated process for revoking system access by disabling accounts immediately upon termination or change of responsibilities of an employee or contractor . Disabling these accounts, instead of deleting accounts, allows preservation of audit trails.		●	●

Internal Only - General

5.3.3.1.2 Ensure password unlock time is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`unlock_time=<n>` - The access will be re-enabled after *<n>* seconds after the lock out. The value `0` has the same meaning as value never - the access will not be re-enabled without resetting the faillock entries by the faillock(8) command.

Note:

- The default directory that pam_faillock uses is usually cleared on system boot so the access will be also re-enabled after system reboot. If that is undesirable a different tally directory must be set with the dir option.
- It is usually undesirable to permanently lock out users as they can become easily a target of denial of service attack unless the usernames are random and kept secret to potential attackers.
- The maximum configurable value for `unlock_time` is `604800`

Rationale:

Locking out user IDs after *n* unsuccessful consecutive login attempts mitigates brute force password attacks against your systems.

Impact:

Use of `unlock_time=0` may allow an attacker to cause denial of service to legitimate users. This will also require a systems administrator with elevated privileges to unlock the account.

Audit:

Run the following command to verify that the time in seconds before the account is unlocked is either `0` (never) or `900` (15 minutes) or more and meets local site policy:

```
# grep -Pi -- '^\h*unlock_time\h*=\h*(0|9[0-9][0-9]|[1-9][0-9]{3,})\b'
/etc/security/faillock.conf

unlock_time = 900
```

Run the following command to verify that the `unlock_time` argument has not been set, or is either `0` (never) or `900` (15 minutes) or more and meets local site policy:

Internal Only - General

```
# grep -Pi --
'^\h*auth\h+(requisite|required|sufficient)\h+pam_faillock\.so\h+([^#\n\r]+\h
+)?unlock_time\h*=\h*([1-9]|[1-9][0-9]|[1-8][0-9][0-9])\b' /etc/pam.d/common-
auth
```

Nothing should be returned

Remediation:

Set password unlock time to conform to site policy. `unlock_time` should be `0` (never), or `900` seconds or greater.

Edit `/etc/security/faillock.conf` and update or add the following line:

```
unlock_time = 900
```

Run the following command: remove the `unlock_time` argument from the `pam_faillock.so` module in the PAM files:

```
# grep -Pl -- '\bpam_faillock\.so\h+([^#\n\r]+\h+)?unlock_time\b'
/usr/share/pam-configs/*
```

Edit any returned files and remove the `unlock_time=<N>` argument from the `pam_faillock.so` line(s).

Default Value:

unlock_time = 600

References:

1. GRID: NIX-00000244

Additional Information:

If a user has been locked out because they have reached the maximum consecutive failure count defined by `deny=` in the `pam_faillock.so` module, the user can be unlocked by issuing the command `faillock --user <USERNAME> --reset`. This command sets the failed count to 0, effectively unlocking the user.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	6.2 Establish an Access Revoking Process Establish and follow a process, preferably automated, for revoking access to enterprise assets, through disabling accounts immediately upon termination, rights revocation, or role change of a user. Disabling accounts, instead of deleting accounts, may be necessary to preserve audit trails.	●	●	●

Internal Only - General

Controls Version	Control	IG 1	IG 2	IG 3
v7	16.7 Establish Process for Revoking Access Establish and follow an automated process for revoking system access by disabling accounts immediately upon termination or change of responsibilities of an employee or contractor . Disabling these accounts, instead of deleting accounts, allows preservation of audit trails.		●	●

Internal Only - General

5.3.3.1.3 Ensure password failed attempts lockout includes root account (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

`even_deny_root` - Enables locking of the root account, as well as regular accounts, after failed attempts. By default, root is exempt to prevent denial-of-service risks.

`root_unlock_time=n` - This option implies even_deny_root option. Allow access after n seconds to root account after the account is locked. In case the option is not specified the value is the same as of the unlock_time option.

Rationale:

Locking out user IDs after `n` unsuccessful consecutive login attempts mitigate brute force password attacks against your systems.

Impact:

Use of `unlock_time=0` or `root_unlock_time=0` may allow an attacker to cause denial of service to legitimate users.

Audit:

Run the following command to verify that `even_deny_root` and/or `root_unlock_time` is enabled:

```
# grep -Pi -- '^\h*(even_deny_root|root_unlock_time\h*=\h*\d+)\b'
/etc/security/faillock.conf
```

Example output:

```
even_deny_root

--AND/OR--

root_unlock_time = 60
```

Run the following command to verify that **- IF -** `root_unlock_time` is set, it is set to `60` (One minute) or more:

```
# grep -Pi -- '^\h*root_unlock_time\h*=\h*([1-9]|[1-5][0-9])\b'
/etc/security/faillock.conf
```

Nothing should be returned

Internal Only - General

Run the following command to check the pam_faillock.so module for the root_unlock_time argument. Verify **-IF-** root_unlock_time is set, it is set to 60 (One minute) or more:

```
# grep -Pi --
'^\h*auth\h+([^#\n\r]+\h+)pam_faillock\.so\h+([^#\n\r]+\h+)?root_unlock_time\
h*=\h*([1-9]|[1-5][0-9])\b' /etc/pam.d/common-auth
```

Nothing should be returned

Remediation:

Edit /etc/security/faillock.conf:

- Remove or update any line containing root_unlock_time, **- OR -** set it to a value of 60 or more
- Update or add the following line

```
even_deny_root
```

Run the following command:

```
# grep -Pl --
'\bpam_faillock\.so\h+([^#\n\r]+\h+)?(even_deny_root|root_unlock_time)'
/usr/share/pam-configs/*
```

Edit any returned files and remove the even_deny_root and root_unlock_time arguments from the pam_faillock.so line(s).

Note: It is recommended to remove any duplicate or contradictory entries of even_deny_root and root_unlock_time found in /usr/share/pam-configs/ files and only use the centralized /etc/security/faillock.conf to set even_deny_root.

Default Value:

disabled

References:

1. GRID: NIX-00000234
2. https://linux.die.net/man/8/pam_faillock

Additional Information:

If a user has been locked out because they have reached the maximum consecutive failure count defined by deny= in the pam_faillock.so module, the user can be unlocked by issuing the command faillock --user <USERNAME> --reset. This command sets the failed count to 0, effectively unlocking the user.

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	6.2 Establish an Access Revoking Process Establish and follow a process, preferably automated, for revoking access to enterprise assets, through disabling accounts immediately upon termination, rights revocation, or role change of a user. Disabling accounts, instead of deleting accounts, may be necessary to preserve audit trails.	●	●	●
v7	16.7 Establish Process for Revoking Access Establish and follow an automated process for revoking system access by disabling accounts immediately upon termination or change of responsibilities of an employee or contractor . Disabling these accounts, instead of deleting accounts, allows preservation of audit trails.		●	●

Internal Only - General

5.3.3.2 Configure pam_pwquality module

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.

These checks are configurable by either:

- use of the module arguments
- modifying the `/etc/security/pwquality.conf` configuration file
- creating a `.conf` file in the `/etc/security/pwquality.conf.d/` directory.

Note: The module arguments override the settings in the `/etc/security/pwquality.conf` configuration file. Settings in the `/etc/security/pwquality.conf` configuration file override settings in a `.conf` file in the `/etc/security/pwquality.conf.d/` directory.

The possible options in the file are:

- `difok` - Number of characters in the new password that must not be present in the old password. (default 1). The special value of 0 disables all checks of similarity of the new password with the old password except the new password being exactly the same as the old one.
- `minlen` - Minimum acceptable size for the new password (plus one if credits are not disabled which is the default). (See pam_pwquality(8).) Cannot be set to lower value than 6. (default 8)
- `dcredit` - The maximum credit for having digits in the new password. If less than 0 it is the minimum number of digits in the new password. (default 0)
- `ucredit` - The maximum credit for having uppercase characters in the new password. If less than 0 it is the minimum number of uppercase characters in the new password. (default 0)
- `lcredit` - The maximum credit for having lowercase characters in the new password. If less than 0 it is the minimum number of lowercase characters in the new password. (default 0)
- `ocredit` - The maximum credit for having other characters in the new password. If less than 0 it is the minimum number of other characters in the new password. (default 0)
- `minclass` - The minimum number of required classes of characters for the new password (digits, uppercase, lowercase, others). (default 0)
- `maxrepeat` - The maximum number of allowed same consecutive characters in the new password. The check is disabled if the value is 0. (default 0)
- `maxsequence` - The maximum length of monotonic character sequences in the new password. Examples of such sequence are '12345' or 'fedcb'. Note that most such passwords will not pass the simplicity check unless the sequence is only a minor part of the password. The check is disabled if the value is 0. (default 0)

Internal Only - General

- `maxclassrepeat` - The maximum number of allowed consecutive characters of the same class in the new password. The check is disabled if the value is 0. (default 0)
- `gecoscheck` - If nonzero, check whether the words longer than 3 characters from the GECOS field of the user's passwd(5) entry are contained in the new password. The check is disabled if the value is 0. (default 0)
- `dictcheck` - If nonzero, check whether the password (with possible modifications) matches a word in a dictionary. Currently the dictionary check is performed

  using the cracklib library. (default 1)

- `usercheck=<N>` - If nonzero, check whether the password (with possible modifications) contains the user name in some form. It is not performed for user names

  shorter than 3 characters. (default 1)

- `usersubstr=<N>` - If greater than 3 (due to the minimum length in usercheck), check whether the password contains a substring of at least N length in some form.

  (default 0)

- `enforcing=<N>` - If nonzero, reject the password if it fails the checks, otherwise only print the warning. This setting applies only to the pam_pwquality module

  and possibly other applications that explicitly change their behavior based on it. It does not affect pwmake(1) and pwscore(1). (default 1)

- `badwords` - Space separated list of words that must not be contained in the password. These are additional words to the cracklib dictionary check. This setting can be also used by applications to emulate the gecos check for user accounts that are not created yet.
- `dictpath` - Path to the cracklib dictionaries. Default is to use the cracklib default.
- `retry=<N>` - Prompt user at most N times before returning with error. The default is 1.
- `enforce_for_root` - The module will return error on failed check even if the user changing the password is root. This option is off by default which means that just the message about the failed check is printed but root can change the password anyway. Note that root is not asked for an old password so the checks that compare the old and new password are not performed.
- `local_users_only` - The module will not test the password quality for users that are not present in the /etc/passwd file. The module still asks for the password so the following modules in the stack can use the use_authtok option. This option is off by default.

Internal Only - General

5.3.3.2.1 Ensure password number of changed characters is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `pwquality difok` option sets the number of characters in a password that must not be present in the old password.

Rationale:

Use of a complex password helps to increase the time and resources required to compromise the password. Password complexity, or strength, is a measure of the effectiveness of a password in resisting attempts at guessing and brute-force attacks.

Password complexity is one factor of several that determines how long it takes to crack a password. The more complex the password, the greater the number of possible combinations that need to be tested before the password is compromised.

Impact:

Requiring a minimum number of changed characters from the previous password reduces the risk of incremental password changes (e.g., changing only one character) that allow users to effectively reuse their previous password.

Audit:

Run the following command to verify that the `difok` option is set to `2` or more and follows local site policy:

```
# grep -Psi -- '^\h*difok\h*=\h*([2-9]|[1-9][0-9]+)\b'
/etc/security/pwquality.conf /etc/security/pwquality.conf.d/*.conf
2>/dev/null
```

Example output:

```
/etc/security/pwquality.conf.d/50-pwdifok.conf:difok = 2
```

Verify returned value(s) are `2` or more and meet local site policy

Run the following command to verify that `difok` is not set, is `2` or more, and conforms to local site policy:

```
# grep -Psi --
'^\h*password\h+(requisite|required|sufficient)\h+pam_pwquality\.so\h+([^#\n\
r]+\h+)?difok\h*=\h*([0-1])\b' /etc/pam.d/common-password
```

Nothing should be returned

Internal Only - General

Note:

- settings should be configured in only **one** location for clarity
- Settings observe an order of precedence:
- module arguments override the settings in the `/etc/security/pwquality.conf` configuration file
- settings in the `/etc/security/pwquality.conf` configuration file override settings in a `.conf` file in the `/etc/security/pwquality.conf.d/` directory
- settings in a `.conf` file in the `/etc/security/pwquality.conf.d/` directory are read in canonical order, with last read file containing the setting taking precedence
- It is recommended that settings be configured in a `.conf` file in the `/etc/security/pwquality.conf.d/` directory for clarity, convenience, and durability.

Remediation:

Create or modify a file ending in `.conf` in the `/etc/security/pwquality.conf.d/` directory or the file `/etc/security/pwquality.conf` and add or modify the following line to set `difok` to `2` or more. Ensure setting conforms to local site policy:

Example:

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

{
   sed -ri 's/^\s*difok\s*=/# &/' /etc/security/pwquality.conf
   [ ! -d /etc/security/pwquality.conf.d/ ] && mkdir
/etc/security/pwquality.conf.d/
   printf '\n%s' "difok = 2" > /etc/security/pwquality.conf.d/50-pwdifok.conf
}
```

Run the following command:

```
# grep -Pl -- '\bpam_pwquality\.so\h+([^#\n\r]+\h+)?difok\b' /usr/share/pam-
configs/*
```

Edit any returned files and remove the `difok` argument from the `pam_pwquality.so` line(s):

Default Value:

difok = 1

References:

1. GRID: NIX-00000240
2. NIST SP 800-53 Rev. 5: IA-5

Internal Only - General

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.		●	●

Internal Only - General

5.3.3.2.2 Ensure password length is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The minimum password length setting determines the lowers number of characters that make up a password for a user account. There are many different theories about how to determine the best password length for an organization, but perhaps "passphrase" is a better term than "password".

The `minlen` option sets the minimum acceptable size for the new password (plus one if credits are not disabled which is the default). Cannot be set to lower value than 6.

Rationale:

Strong passwords help protect systems from password attacks. Types of password attacks include dictionary attacks, which attempt to use common words and phrases, and brute force attacks, which try every possible combination of characters. Also attackers may try to obtain the account database so they can use tools to discover the accounts and passwords.

Impact:

In general, it is true that longer passwords are better (harder to crack), but it is also true that forced password length requirements can cause user behavior that is predictable and undesirable. For example, requiring users to have a minimum 16-character password may cause them to choose repeating patterns like fourfourfourfour or passwordpassword that meet the requirement but aren't hard to guess. Additionally, length requirements increase the chances that users will adopt other insecure practices, like writing them down, re-using them or storing them unencrypted in their documents.

Having a reasonable minimum length with no maximum character limit increases the resulting average password length used (and therefore the strength).6

Audit:

Run the following command to verify that password length is 14 or more characters, and conforms to local site policy:

```
# grep -Psi -- '^\h*minlen\h*=\h*(1[4-9]|[2-9][0-9]|[1-9][0-9]{2,})\b'
/etc/security/pwquality.conf /etc/security/pwquality.conf.d/*.conf
2>/dev/null
```

Example output:

Internal Only - General

```
/etc/security/pwquality.conf.d/50-pwlength.conf:minlen = 14
```

Verify returned value(s) are no less than 14 characters and meet local site policy

Run the following command to verify that minlen is not set, or is 14 or more characters, and conforms to local site policy:

```
# grep -Psi --
'^\h*password\h+(requisite|required|sufficient)\h+pam_pwquality\.so\h+([^#\n\
r]+\h+)?minlen\h*=\h*([0-9]|1[0-3])\b' /etc/pam.d/common-password
```

Nothing should be returned

Note:

- settings should be configured in only **one** location for clarity
- Settings observe an order of precedence:
- module arguments override the settings in the /etc/security/pwquality.conf configuration file
- settings in the /etc/security/pwquality.conf configuration file override settings in a .conf file in the /etc/security/pwquality.conf.d/ directory
- settings in a .conf file in the /etc/security/pwquality.conf.d/ directory are read in canonical order, with last read file containing the setting taking precedence
- It is recommended that settings be configured in a .conf file in the /etc/security/pwquality.conf.d/ directory for clarity, convenience, and durability.

Remediation:

Create or modify a file ending in .conf in the /etc/security/pwquality.conf.d/ directory or the file /etc/security/pwquality.conf and add or modify the following line to set password length of 14 or more characters. Ensure that password length conforms to local site policy:

Example:

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

{
   sed -ri 's/^\s*minlen\s*=/# &/' /etc/security/pwquality.conf
   [ ! -d /etc/security/pwquality.conf.d/ ] && mkdir
/etc/security/pwquality.conf.d/
   printf '\n%s' "minlen = 14" > /etc/security/pwquality.conf.d/50-
pwlength.conf
}
```

Run the following command:

Internal Only - General

```
# grep -Pl -- '\bpam_pwquality\.so\h+([^#\n\r]+\h+)?minlen\b' /usr/share/pam-configs/*
```

Edit any returned files and remove the `minlen` argument from the `pam_pwquality.so` line(s):

Default Value:

minlen = 8

References:

1. GRID: NIX-00000238
2. CIS Password Policy Guide
3. NIST SP 800-53 Rev. 5: IA-5(1)
4. pam_pwquality(8)

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.		●	●

Internal Only - General

5.3.3.2.3 Ensure password complexity is configured (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Password complexity can be set through:

- `minclass` - The minimum number of classes of characters required in a new password. (digits, uppercase, lowercase, others). e.g. `minclass = 4` requires digits, uppercase, lower case, and special characters.
- `dcredit` - The maximum credit for having digits in the new password. If less than `0` it is the minimum number of digits in the new password. e.g. `dcredit = -1` requires at least one digit
- `ucredit` - The maximum credit for having uppercase characters in the new password. If less than `0` it is the minimum number of uppercase characters in the new password. e.g. `ucredit = -1` requires at least one uppercase character
- `ocredit` - The maximum credit for having other characters in the new password. If less than `0` it is the minimum number of other characters in the new password. e.g. `ocredit = -1` requires at least one special character
- `lcredit` - The maximum credit for having lowercase characters in the new password. If less than `0` it is the minimum number of lowercase characters in the new password. e.g. `lcredit = -1` requires at least one lowercase character

Rationale:

Strong passwords protect systems from being hacked through brute force methods.

Requiring at least one non-alphabetic character increases the search space beyond pure dictionary words, which makes the resulting password harder to crack.

Forcing users to choose an excessively complex password, e.g. some combination of upper-case, lower-case, numbers, and special characters, has a negative impact. It places an extra burden on users and many will use predictable patterns (for example, a capital letter in the first position, followed by lowercase letters, then one or two numbers, and a “special character” at the end). Attackers know this, so dictionary attacks will often contain these common patterns and use the most common substitutions like, $ for s, @ for a, 1 for l, 0 for o.

Impact:

CIS password complexity requirements may differ from other frameworks or policies, adherence to site-specific policy is imperative.

Internal Only - General

Passwords that are too complex in nature make it harder for users to remember, leading to bad practices. In addition, composition requirements provide no defense against common attack types such as social engineering or insecure storage of passwords.

Audit:

Run the following command to verify:

- `dcredit`, `ucredit`, `lcredit`, and `ocredit` are not set to a value greater than `0`
- Complexity conforms to local site policy:

```
# grep -Psi -- '^\h*(minclass|[dulo]credit)\b' /etc/security/pwquality.conf
/etc/security/pwquality.conf.d/*.conf 2>/dev/null
```

Example output:

```
/etc/security/pwquality.conf.d/50-pwcomplexity.conf:minclass = 3
/etc/security/pwquality.conf.d/50-pwcomplexity.conf:ucredit = -2
/etc/security/pwquality.conf.d/50-pwcomplexity.conf:lcredit = -2
/etc/security/pwquality.conf.d/50-pwcomplexity.conf:dcredit = -1
/etc/security/pwquality.conf.d/50-pwcomplexity.conf:ocredit = 0
```

The example represents a requirement of three character classes, with passwords requiring two upper case, two lower case, and one numeric character.

Run the following command to verify that module arguments in the configuration file(s) are not being overridden by arguments in `/etc/pam.d/common-password`:

```
# grep -Psi --
'^\h*password\h+(requisite|required|sufficient)\h+pam_pwquality\.so\h+([^#\n\
r]+\h+)?(minclass=\d*|[dulo]credit=-?\d*)\b' /etc/pam.d/common-password
```

Nothing should be returned

Note:

- settings should be configured in only **one** location for clarity
- Settings observe an order of precedence:
- module arguments override the settings in the `/etc/security/pwquality.conf` configuration file
- settings in the `/etc/security/pwquality.conf` configuration file override settings in a `.conf` file in the `/etc/security/pwquality.conf.d/` directory
- settings in a `.conf` file in the `/etc/security/pwquality.conf.d/` directory are read in canonical order, with last read file containing the setting taking precedence
- It is recommended that settings be configured in a `.conf` file in the `/etc/security/pwquality.conf.d/` directory for clarity, convenience, and durability.

Internal Only - General

Remediation:

Note: CIS password complexity requirements may differ from other frameworks or policies, adherence to site-specific policy is imperative.

Run the following command:

```
# grep -Pl --
'\bpam_pwquality\.so\h+([^#\n\r]+\h+)?(minclass|[dulo]credit)\b'
/usr/share/pam-configs/*
```

Edit any returned files and remove the `minclass`, `dcredit`, `ucredit`, `lcredit`, and `ocredit` arguments from the `pam_pwquality.so` line(s)

Create or modify a file ending in `.conf` in the `/etc/security/pwquality.conf.d/` directory or the file `/etc/security/pwquality.conf` and add or modify the following line(s) to set complexity according to local site policy:

- `minclass = _N_`
- `dcredit = _N_` # Value should be either `0` or a number proceeded by a minus (`-`) symbol
- `ucredit = -1` # Value should be either `0` or a number proceeded by a minus (`-`) symbol
- `ocredit = -1` # Value should be either `0` or a number proceeded by a minus (`-`) symbol
- `lcredit = -1` # Value should be either `0` or a number proceeded by a minus (`-`) symbol

Example 1 - Set `minclass = 3`:

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

{
   sed -ri 's/^\s*minclass\s*=/# &/' /etc/security/pwquality.conf
   sed -ri 's/^\s*[dulo]credit\s*=/# &/' /etc/security/pwquality.conf
   [ ! -d /etc/security/pwquality.conf.d/ ] && mkdir
/etc/security/pwquality.conf.d/
   printf '\n%s' "minclass = 3" > /etc/security/pwquality.conf.d/50-
pwcomplexity.conf
}
```

Example 2 - set `dcredit = -1`, `ucredit = -1`, and `lcredit = -1`:

Internal Only - General

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

{
   sed -ri 's/^\s*minclass\s*=/# &/' /etc/security/pwquality.conf
   sed -ri 's/^\s*[dulo]credit\s*=/# &/' /etc/security/pwquality.conf
   [ ! -d /etc/security/pwquality.conf.d/ ] && mkdir
/etc/security/pwquality.conf.d/
   printf '%s\n' "dcredit = -1" "ucredit = -1" "lcredit = -1" >
/etc/security/pwquality.conf.d/50-pwcomplexity.conf
}
```

Default Value:

minclass = 0

dcredit = 0

ucredit = 0

ocredit = 0

lcredit = 0

References:

1. GRID: NIX-00000232
2. PWQUALITY.CONF(5)
3. https://www.cisecurity.org/insights/white-papers/cis-password-policy-guide
4. NIST SP 800-53 Rev. 5: IA-5
5. pam_pwquality(8)

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.		●	●

Internal Only - General

5.3.3.2.4 Ensure password same consecutive characters is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `pwquality maxrepeat` option sets the maximum number of allowed same consecutive characters in a new password.

Rationale:

Use of a complex password helps to increase the time and resources required to compromise the password. Password complexity, or strength, is a measure of the effectiveness of a password in resisting attempts at guessing and brute-force attacks.

Password complexity is one factor of several that determines how long it takes to crack a password. The more complex the password, the greater the number of possible combinations that need to be tested before the password is compromised.

Impact:

Limiting consecutive repeated characters prevents trivially simple passwords such as 'aaabbbccc' or '111222333' that technically meet length requirements but offer little resistance to pattern-based attacks.

Audit:

Run the following command to verify that the `maxrepeat` option is set to `3` or less, not `0`, and follows local site policy:

```
# grep -Psi -- '^\h*maxrepeat\h*=\h*[1-3]\b' /etc/security/pwquality.conf
/etc/security/pwquality.conf.d/*.conf 2>/dev/null
```

Example output:

```
/etc/security/pwquality.conf.d/50-pwrepeat.conf:maxrepeat = 3
```

Verify returned value(s) are `3` or less, not `0`, and meet local site policy

Run the following command to verify that `maxrepeat` is not set, is `3` or less, not `0`, and conforms to local site policy:

```
# grep -Psi --
'^\h*password\h+(requisite|required|sufficient)\h+pam_pwquality\.so\h+([^#\n\
r]+\h+)?maxrepeat\h*=\h*(0|[4-9]|[1-9][0-9]+)\b' /etc/pam.d/common-password
```

Nothing should be returned

Internal Only - General

Note:

- settings should be configured in only **one** location for clarity
- Settings observe an order of precedence:
- module arguments override the settings in the `/etc/security/pwquality.conf` configuration file
- settings in the `/etc/security/pwquality.conf` configuration file override settings in a `.conf` file in the `/etc/security/pwquality.conf.d/` directory
- settings in a `.conf` file in the `/etc/security/pwquality.conf.d/` directory are read in canonical order, with last read file containing the setting taking precedence
- It is recommended that settings be configured in a `.conf` file in the `/etc/security/pwquality.conf.d/` directory for clarity, convenience, and durability.

Remediation:

Create or modify a file ending in `.conf` in the `/etc/security/pwquality.conf.d/` directory or the file `/etc/security/pwquality.conf` and add or modify the following line to set `maxrepeat` to `3` or less and not `0`. Ensure setting conforms to local site policy:

Example:

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

{
 sed -ri 's/^\s*maxrepeat\s*=/# &/' /etc/security/pwquality.conf
 [ ! -d /etc/security/pwquality.conf.d/ ] && mkdir
/etc/security/pwquality.conf.d/
 printf '\n%s' "maxrepeat = 3" > /etc/security/pwquality.conf.d/50-
pwrepeat.conf
}
```

Run the following command to check if any pam-configs profile sets `maxrepeat`:

```
# grep -Pl -- '\bpam_pwquality\.so\h+([^#\n\r]+\h+)?maxrepeat\b'
/usr/share/pam-configs/*
```

Edit any returned files and remove the `maxrepeat` argument from the `pam_pwquality.so` line(s).

Default Value:

maxrepeat = 0

References:

1. GRID: NIX-00000243
2. NIST SP 800-53 Rev. 5: IA-5

Internal Only - General

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.		●	●

Internal Only - General

5.3.3.2.5 Ensure password maximum sequential characters is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `pwquality maxsequence` option sets the maximum length of monotonic character sequences in the new password. Examples of such sequence are `12345` or `fedcb`. The check is disabled if the value is `0`.

Note: Most such passwords will not pass the simplicity check unless the sequence is only a minor part of the password.

Rationale:

Use of a complex password helps to increase the time and resources required to compromise the password. Password complexity, or strength, is a measure of the effectiveness of a password in resisting attempts at guessing and brute-force attacks.

Password complexity is one factor of several that determines how long it takes to crack a password. The more complex the password, the greater the number of possible combinations that need to be tested before the password is compromised.

Impact:

Limiting sequential character runs (e.g., '12345', 'fedcb') prevents passwords that appear complex in length but are trivially guessable via ordered-sequence dictionaries.

Audit:

Run the following command to verify that the `maxsequence` option is set to `3` or less, not `0`, and follows local site policy:

```
# grep -Psi -- '^\h*maxsequence\h*=\h*[1-3]\b' /etc/security/pwquality.conf
/etc/security/pwquality.conf.d/*.conf 2>/dev/null
```

Example output:

```
/etc/security/pwquality.conf.d/50-pwmaxsequence.conf:maxsequence = 3
```

Verify returned value(s) are `3` or less, not `0`, and meet local site policy

Run the following command to verify that `maxsequence` is not set, is `3` or less, not `0`, and conforms to local site policy:

Internal Only - General

```
# grep -Psi --
'^\h*password\h+(requisite|required|sufficient)\h+pam_pwquality\.so\h+([^#\n\
r]+\h+)?maxsequence\h*=\h*(0|[4-9]|[1-9][0-9]+)\b' /etc/pam.d/common-password
```

Nothing should be returned

Note:

- settings should be configured in only **one** location for clarity
- Settings observe an order of precedence:
- module arguments override the settings in the `/etc/security/pwquality.conf` configuration file
- settings in the `/etc/security/pwquality.conf` configuration file override settings in a `.conf` file in the `/etc/security/pwquality.conf.d/` directory
- settings in a `.conf` file in the `/etc/security/pwquality.conf.d/` directory are read in canonical order, with last read file containing the setting taking precedence
- It is recommended that settings be configured in a `.conf` file in the `/etc/security/pwquality.conf.d/` directory for clarity, convenience, and durability.

Remediation:

Create or modify a file ending in `.conf` in the `/etc/security/pwquality.conf.d/` directory or the file `/etc/security/pwquality.conf` and add or modify the following line to set `maxsequence` to `3` or less and not `0`. Ensure setting conforms to local site policy:

Example:

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

{
   sed -ri 's/^\s*maxsequence\s*=/# &/' /etc/security/pwquality.conf
   [ ! -d /etc/security/pwquality.conf.d/ ] && mkdir
/etc/security/pwquality.conf.d/
   printf '\n%s' "maxsequence = 3" > /etc/security/pwquality.conf.d/50-
pwmaxsequence.conf
}
```

Run the following command:

```
# grep -Pl -- '\bpam_pwquality\.so\h+([^#\n\r]+\h+)?maxsequence\b'
/usr/share/pam-configs/*
```

Edit any returned files and remove the `maxsequence` argument from the `pam_pwquality.so` line(s):

Default Value:

maxsequence = 0

Internal Only - General

References:

1. GRID: NIX-00000239
2. NIST SP 800-53 Rev. 5: IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.2 Use Unique Passwords Use unique passwords for all enterprise assets. Best practice implementation includes, at a minimum, an 8-character password for accounts using MFA and a 14-character password for accounts not using MFA.	●	●	●
v7	4.4 Use Unique Passwords Where multi-factor authentication is not supported (such as local administrator, root, or service accounts), accounts will use passwords that are unique to that system.		●	●

Internal Only - General

5.3.3.2.6 Ensure password dictionary check is enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `pwquality dictcheck` option sets whether to check for the words from the `cracklib` dictionary.

Rationale:

If the operating system allows the user to select passwords based on dictionary words, this increases the chances of password compromise by increasing the opportunity for successful guesses, and brute-force attacks.

Impact:

Without dictionary checking, users can set passwords that are common words or simple modifications of words, which are trivially crackable via dictionary attacks even when they meet length and complexity requirements.

Audit:

Run the following command to verify that the `dictcheck` option is not set to `0` (disabled) in a pwquality configuration file:

```
# grep -Psi -- '^\h*dictcheck\h*=\h*0\b' /etc/security/pwquality.conf
/etc/security/pwquality.conf.d/*.conf 2>/dev/null
```

Nothing should be returned

Run the following command to verify that the `dictcheck` option is not set to `0` (disabled) as a module argument in a PAM file:

```
# grep -Psi --
'^\h*password\h+(requisite|required|sufficient)\h+pam_pwquality\.so\h+([^#\n\
r]+\h+)?dictcheck\h*=\h*0\b' /etc/pam.d/common-password
```

Nothing should be returned

Note:

- Settings observe an order of precedence:
- module arguments override the settings in the `/etc/security/pwquality.conf` configuration file
- settings in the `/etc/security/pwquality.conf` configuration file override settings in a `.conf` file in the `/etc/security/pwquality.conf.d/` directory

Internal Only - General

- settings in a `.conf` file in the `/etc/security/pwquality.conf.d/` directory are read in canonical order, with last read file containing the setting taking precedence
- It is recommended that settings be configured in a `.conf` file in the `/etc/security/pwquality.conf.d/` directory for clarity, convenience, and durability.

Remediation:

Edit any file ending in `.conf` in the `/etc/security/pwquality.conf.d/` directory and/or the file `/etc/security/pwquality.conf` and comment out or remove any instance of `dictcheck = 0`:

Example:

```
# sed -ri 's/^\s*dictcheck\s*=/# &/' /etc/security/pwquality.conf
/etc/security/pwquality.conf.d/*.conf 2>/dev/null
```

Run the following command:

```
# grep -Pl -- '\bpam_pwquality\.so\h+([^#\n\r]+\h+)?dictcheck\b'
/usr/share/pam-configs/*
```

Edit any returned files and remove the `dictcheck` argument from the `pam_pwquality.so` line(s)

Default Value:

dictcheck = 1

References:

1. GRID: NIX-00000233
2. NIST SP 800-53 Rev. 5: IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.2 Use Unique Passwords Use unique passwords for all enterprise assets. Best practice implementation includes, at a minimum, an 8-character password for accounts using MFA and a 14-character password for accounts not using MFA.	●	●	●
v7	4.4 Use Unique Passwords Where multi-factor authentication is not supported (such as local administrator, root, or service accounts), accounts will use passwords that are unique to that system.		●	●

Internal Only - General

5.3.3.2.7 Ensure password quality checking is enforced (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `pam_pwquality` module can be configured to either reject a password if it fails the checks, or only print a warning.

This is configured by setting the `enforcing=<N>` argument. If nonzero, a password will be rejected if it fails the checks, otherwise only a warning message will be provided.

This setting applies only to the pam_pwquality module and possibly other applications that explicitly change their behavior based on it. It does not affect pwmake(1) and pwscore(1).

Rationale:

Strong passwords help protect systems from password attacks. Types of password attacks include dictionary attacks, which attempt to use common words and phrases, and brute force attacks, which try every possible combination of characters. Also attackers may try to obtain the account database so they can use tools to discover the accounts and passwords.

Impact:

Setting `enforcing=0` downgrades `pam_pwquality` to advisory mode — password quality failures are logged but the password is accepted anyway. This renders all other password quality settings (5.3.3.2.1–.6, .8) ineffective.

Audit:

Run the following command to verify that `enforcing=0` has not been set in a `pwquality` configuration file:

```
# grep -PHsi -- '^\h*enforcing\h*=\h*0\b' /etc/security/pwquality.conf
/etc/security/pwquality.conf.d/*.conf 2>/dev/null
```

Nothing should be returned

Run the following command to verify that the `enforcing=0` argument has not been set on the `pam_pwquality` module:

Internal Only - General

```
# grep -PHsi --
'^\h*password\h+[^#\n\r]+\h+pam_pwquality\.so\h+([^#\n\r]+\h+)?enforcing=0\b'
/etc/pam.d/common-password
```

Nothing should be returned

Remediation:

Run the following command:

```
# grep -Pl -- '\bpam_pwquality\.so\h+([^#\n\r]+\h+)?enforcing=0\b'
/usr/share/pam-configs/*
```

Edit any returned files and remove the `enforcing=0` argument from the `pam_pwquality.so` line(s)

Edit `/etc/security/pwquality.conf` and all files ending in `.conf` in the `/etc/security/pwquality.conf.d/` directory and remove or comment out any line containing the `enforcing = 0` argument:

Example:

```
# sed -ri 's/^\s*enforcing\s*=\s*0/# &/' /etc/security/pwquality.conf
/etc/security/pwquality.conf.d/*.conf
```

Default Value:

enforcing=1

References:

1. GRID: NIX-00000241
2. PWQUALITY.CONF(5)
3. pam_pwquality(8)
4. NIST SP 800-53 Rev. 5: IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.2 Use Unique Passwords Use unique passwords for all enterprise assets. Best practice implementation includes, at a minimum, an 8-character password for accounts using MFA and a 14-character password for accounts not using MFA.	●	●	●
v7	4.4 Use Unique Passwords Where multi-factor authentication is not supported (such as local administrator, root, or service accounts), accounts will use passwords that are unique to that system.		●	●

Internal Only - General

5.3.3.2.8 Ensure password quality is enforced for the root user (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

If the `pwquality enforce_for_root` option is enabled, the module will return error on failed check even if the user changing the password is root.

This option is off by default which means that just the message about the failed check is printed but root can change the password anyway.

Note: The root is not asked for an old password so the checks that compare the old and new password are not performed.

Rationale:

Use of a complex password helps to increase the time and resources required to compromise the password. Password complexity, or strength, is a measure of the effectiveness of a password in resisting attempts at guessing and brute-force attacks.

Password complexity is one factor of several that determines how long it takes to crack a password. The more complex the password, the greater the number of possible combinations that need to be tested before the password is compromised.

Impact:

Without `enforce_for_root`, the root user can bypass all `pam_pwquality` password strength checks. Root-set passwords that fail quality checks will still be accepted with only a warning. Enabling this option ensures root passwords meet the same standards as regular users.

Audit:

Run the following command to verify that the `enforce_for_root` option is enabled in a pwquality configuration file:

```
# grep -Psi -- '^\h*enforce_for_root\b' /etc/security/pwquality.conf
/etc/security/pwquality.conf.d/*.conf 2>/dev/null
```

Example output:

```
/etc/security/pwquality.conf.d/50-pwroot.conf:enforce_for_root
```

Note:

- Settings observe an order of precedence:

Internal Only - General

- module arguments override the settings in the `/etc/security/pwquality.conf` configuration file
- settings in the `/etc/security/pwquality.conf` configuration file override settings in a `.conf` file in the `/etc/security/pwquality.conf.d/` directory
- settings in a `.conf` file in the `/etc/security/pwquality.conf.d/` directory are read in canonical order, with last read file containing the setting taking precedence
- It is recommended that settings be configured in a `.conf` file in the `/etc/security/pwquality.conf.d/` directory for clarity, convenience, and durability.

Remediation:

Edit or add the following line in a `*.conf` file in `/etc/security/pwquality.conf.d` or in `/etc/security/pwquality.conf`:

Example:

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

{
   [ ! -d /etc/security/pwquality.conf.d/ ] && mkdir
/etc/security/pwquality.conf.d/
   printf '\n%s\n' "enforce_for_root" > /etc/security/pwquality.conf.d/50-
pwroot.conf
}
```

Default Value:

disabled

References:

1. GRID: NIX-00000242
2. NIST SP 800-53 Rev. 5: IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.2 Use Unique Passwords Use unique passwords for all enterprise assets. Best practice implementation includes, at a minimum, an 8-character password for accounts using MFA and a 14-character password for accounts not using MFA.	●	●	●
v7	4.4 Use Unique Passwords Where multi-factor authentication is not supported (such as local administrator, root, or service accounts), accounts will use passwords that are unique to that system.		●	●

Internal Only - General

Internal Only - General

5.3.3.3 Configure pam_pwhistory module

`pam_history` module - This PAM module saves the last passwords for each user in order to force password change history and keep the user from alternating between the same password too frequently.

This module does not work together with kerberos. In general, it does not make much sense to use this module in conjunction with `NIS` or `LDAP`, since the old passwords are stored on the local machine and are not available on another machine for password history checking

`pwhistory.conf` provides a way to configure the default settings for saving the last passwords for each user. This file is read by the `pam_pwhistory` module and is the preferred method over configuring `pam_pwhistory` directly.

The file has a very simple `name = value` format with possible comments starting with `#` character. The whitespace at the beginning of line, end of line, and around the `=` sign is ignored.

Options:

- `debug` - Turns on debugging via syslog(3).
- `use_authtok` - When password changing enforce the module to use the new password provided by a previously stacked password module (this is used in the example of the stacking of the `pam_passwdqc module` documented below).
- `enforce_for_root` - If this option is set, the check is enforced for root, too.
- `remember=<N>` - The last <*N*> passwords for each user are saved. The default is `10`. Value of `0` makes the module to keep the existing contents of the opasswd file unchanged.
- `retry=<N>` - Prompt user at most <*N*> times before returning with error. The default is `1`.
- `authtok_type=<STRING>` - See pam_get_authtok(3) for more details.

Examples:

`/etc/security/pwhistory.conf` file example:

```
debug
remember=5
file=/tmp/opasswd
```

Note: The options specified on the `pwhistory.so` module command line override the values from the `pwhistory.conf` configuration file. The `pwhistory.conf` file is the preferred method over configuring `pam_pwhistory` directly.

Internal Only - General

5.3.3.3.1 Ensure password history remember is configured (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. The number of passwords remembered is set via the remember argument value in set for the `pam_pwhistory` module.

The `pwhistory.conf` provides a way to configure the default settings for saving the last passwords for each user. This file is read by the `pam_pwhistory` module and is the preferred method over configuring `pam_pwhistory` directly.

- `remember=<N>` - The last `<N>` passwords for each user are saved. The default is `10`. Value of `0` makes the module to keep the existing contents of the `opasswd` file unchanged.

Rationale:

Requiring users not to reuse their passwords make it less likely that an attacker will be able to guess the password or use a compromised password.

> **Note:** These changes only apply to accounts configured on the local system.

Audit:

Run the following command and verify that the remember option is set to `24` or more and meets local site policy in `/etc/security/pwhistory.conf`:

```
# grep -Pi -- '^\h*remember\h*=\h*(2[4-9]|[3-9][0-9]|[1-9][0-9]{2,})\b'
/etc/security/pwhistory.conf

remember = 24
```

- OR/IF - The `pam_pwhistory` profile in `/usr/share/pam-configs/` has been used to configure `pwhistory`

Run the following command and verify:

- The `pwhistory` line in `/etc/pam.d/common-password` includes remember=<N>
- The value of `<N>` is `24` or more
- The value meets local site policy

Internal Only - General

```
# grep -P -- '\bpam_pwhistory\.so\b' /etc/pam.d/common-password
```

Output should be similar to:

```
password   requisite   pam_pwhistory.so remember=24 enforce_for_root
use_authtok
```

Note:

- The options specified on the `pwhistory.so` module command line override the values from the `pwhistory.conf` configuration file. The `pwhistory.conf` file is the preferred method over configuring `pam_pwhistory` directly.
- Only one method should be used to set the values for `pwhistory`. If output is provided for both locations this is considered a failure.

Remediation:

Note: The options specified on the `pam_pwhistory.so` module command line override the values from the `pwhistory.conf` configuration file. The `pwhistory.conf` file is the preferred method over configuring `pam_pwhistory` directly. Only one method should be used; if `remember=<N>` is set in both locations and the values differ, the per-module argument wins and the `pwhistory.conf` value is silently ignored.

Step 1 - Set the value in `pwhistory.conf` (preferred method):

Edit `/etc/security/pwhistory.conf` and add or modify the `remember` option, setting it to `24` or more depending on local site policy:

```
remember = 24
```

Step 2 (conditional) - Reconcile any conflicting `pam-configs` profile:

If a profile under `/usr/share/pam-configs/` sets `remember=<N>` on the `pam_pwhistory.so` line and the value differs from what was set in Step 1, the per-module argument will override `pwhistory.conf`. List the candidate profiles:

```
# grep -RPl '\bpam_pwhistory\.so\b' /usr/share/pam-configs/
```

For each profile returned, open the file in a text editor and locate the `Password:` and `Password-Initial:` blocks. Remove the `remember=<N>` argument from the `pam_pwhistory.so` line in those blocks (leave any other arguments such as `enforce_for_root` or `use_authtok` in place).

Re-render `/etc/pam.d/common-password` from the updated profile by running:

```
# pam-auth-update --enable <profile-name>
```

where `<profile-name>` matches the filename in `/usr/share/pam-configs/` (without any path).

Verification:

Internal Only - General

Confirm that `pwhistory.conf` carries the intended value and that no conflicting per-module argument remains:

```
# grep -Pi -- '^\h*remember\h*=\h*\d+\b' /etc/security/pwhistory.conf
# grep -PHs -- '\bpam_pwhistory\.so\b[^#\n\r]*\bremember=' /etc/pam.d/common-password
```

The first command should return `remember = 24` (or your site-policy value). The second command should return no output.

References:

1. GRID: NIX-00000237
2. https://manpages.debian.org/trixie/libpam-runtime/pwhistory.conf.5.en.html
3. https://manpages.debian.org/trixie/libpam-runtime/pam_pwhistory.8.en.html
4. NIST SP 800-53 Rev. 5: IA-5(1)

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.		●	●

Internal Only - General

5.3.3.3.2 Ensure password history is enforced for the root user (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

If the `pwhistory enforce_for_root` option is enabled, the module will enforce password history for the root user as well

Rationale:

Requiring users not to reuse their passwords make it less likely that an attacker will be able to guess the password or use a compromised password

> **Note:** These changes only apply to accounts configured on the local system.

Audit:

Run the following command to verify that the `enforce_for_root` option is enabled in `/etc/security/pwhistory.conf`:

```
# grep -Pi -- '^\h*enforce_for_root\b' /etc/security/pwhistory.conf

enforce_for_root
```

- OR/IF - The `pam_pwhistory` profile in `/usr/share/pam-configs/` has been used to configure `pwhistory`

Run the following command to verify that the `enforce_for_root` argument exists on the `pwhistory` line in `/etc/pam.d/common-password`:

```
# grep -Psi --
'^\h*password\h+[^#\n\r]+\h+pam_pwhistory\.so\h+([^#\n\r]+\h+)?enforce_for_ro
ot\b' /etc/pam.d/common-password
```

Output should be similar to:

```
password   requisite   pam_pwhistory.so remember=24 enforce_for_root
use_authtok
```

> **Note**:
>
> - The options specified on the `pwhistory.so` module command line override the values from the `pwhistory.conf` configuration file. The `pwhistory.conf` file is the preferred method over configuring `pam_pwhistory` directly.

Internal Only - General

- Only one method should be used to set the values for pwhistory. If output is provided for both locations this is considered a failure.

Remediation:

> **Note:** The options specified on the pam_pwhistory.so module command line override the values from the pwhistory.conf configuration file. The pwhistory.conf file is the preferred method over configuring pam_pwhistory directly. Additionally, if both methods are used to set the same option, pwhistory may silently ignore one — only one method should be used.

Step 1 — Configure /etc/security/pwhistory.conf (preferred method):

Edit /etc/security/pwhistory.conf and add or set the enforce_for_root option:

```
enforce_for_root
```

Step 2 — Reconcile the pam-configs profile (conditional cleanup):

- IF - the pam_pwhistory profile in /usr/share/pam-configs/ has been used to configure pwhistory (i.e. an enforce_for_root argument is present on the pam_pwhistory.so line), remove that argument so the value in pwhistory.conf is the single source of truth. Run the following script:

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

{
profiles=$(grep -Rl "pam_pwhistory.so" /usr/share/pam-configs/ || true)
for profile in $profiles; do
 if grep -Pq '\h+pam_pwhistory\.so\h+([^#\n\r]+\h+)?enforce_for_root\b'
"$profile"; then
 sed -Ei '/pam_pwhistory\.so/s/\benforce_for_root\b//g' "$profile"
 sed -Ei 's/[[:space:]]+/ /g' "$profile"
 echo "Updated pam profile: $profile"
 pam-auth-update --package < /dev/null
 fi
done
}
```

Verification:

After applying Step 1 (and Step 2 if applicable), confirm the setting is in effect:

```
# grep -Pi -- '^\h*enforce_for_root\b' /etc/security/pwhistory.conf
# grep -P -- '\bpam_pwhistory\.so\b' /etc/pam.d/common-password
```

The first command should return enforce_for_root. The second command should return the pam_pwhistory.so line **without** an enforce_for_root argument (so the pwhistory.conf value applies).

Default Value:

disabled

Internal Only - General

References:

1. GRID: NIX-00000236
2. NIST SP 800-53 Rev. 5: IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.2 Use Unique Passwords Use unique passwords for all enterprise assets. Best practice implementation includes, at a minimum, an 8-character password for accounts using MFA and a 14-character password for accounts not using MFA.	●	●	●
v7	4.4 Use Unique Passwords Where multi-factor authentication is not supported (such as local administrator, root, or service accounts), accounts will use passwords that are unique to that system.		●	●

Internal Only - General

5.3.3.3.3 Ensure pam_pwhistory includes use_authtok (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`use_authtok` - When password changing enforce the module to set the new password to the one provided by a previously stacked password module

Rationale:

`use_authtok` allows multiple pam modules to confirm a new password before it is accepted.

Audit:

Run the following command to verify that the `use_authtok` option is enabled in `/etc/security/pwhistory.conf`:

```
# grep -Pi -- '^\h*use_authtok\b' /etc/security/pwhistory.conf

use_authtok
```

- OR/IF - The `pam_pwhistory` profile in `/usr/share/pam-configs/` has been used to configure `pwhistory`

Run the following command to verify that the `use_authtok` argument exists on the pwhistory line in `/etc/pam.d/common-password`:

```
# grep -Psi --
'^\h*password\h+[^#\n\r]+\h+pam_pwhistory\.so\h+([^#\n\r]+\h+)?use_authtok\b'
/etc/pam.d/common-password
```

Output should be similar to:

```
password   requisite   pam_pwhistory.so remember=24 enforce_for_root
use_authtok
```

Note:

- The options specified on the `pwhistory.so` module command line override the values from the `pwhistory.conf` configuration file. The `pwhistory.conf` file is the preferred method over configuring `pam_pwhistory` directly.
- Only one method should be used to set the values for `pwhistory`. If output is provided for both locations this is considered a failure.

Internal Only - General

Remediation:

> **Note:** The options specified on the pam_pwhistory.so module command line override the values from the pwhistory.conf configuration file. The pwhistory.conf file is the preferred method over configuring pam_pwhistory directly. Additionally, if both methods are used to set the same option, pwhistory may silently ignore one — only one method should be used.

Step 1 — Configure /etc/security/pwhistory.conf (preferred method):

Edit /etc/security/pwhistory.conf and add or set the use_authtok option:

```
use_authtok
```

Step 2 — Reconcile the pam-configs profile (conditional cleanup):

- IF - the pam_pwhistory profile in /usr/share/pam-configs/ has been used to configure pwhistory (i.e. a use_authtok argument is present on the pam_pwhistory.so line), remove that argument so the value in pwhistory.conf is the single source of truth. Run the following script:

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

{
profiles=$(grep -Rl "pam_pwhistory.so" /usr/share/pam-configs/ || true)
for profile in $profiles; do
 if grep -Pq '\h+pam_pwhistory\.so\h+([^#\n\r]+\h+)?use_authtok\b'
"$profile"; then
 sed -Ei '/pam_pwhistory\.so/s/\buse_authtok\b//g' "$profile"
 sed -Ei 's/[[:space:]]+/ /g' "$profile"
 echo "Updated pam profile: $profile"
 pam-auth-update --package < /dev/null
 fi
done
}
```

Verification:

After applying Step 1 (and Step 2 if applicable), confirm the setting is in effect:

```
# grep -Pi -- '^\h*use_authtok\b' /etc/security/pwhistory.conf
# grep -P -- '\bpam_pwhistory\.so\b' /etc/pam.d/common-password
```

The first command should return use_authtok. The second command should return the pam_pwhistory.so line **without** a use_authtok argument (so the pwhistory.conf value applies).

References:

1. GRID: NIX-00000224
2. NIST SP 800-53 Rev. 5: IA-5

Internal Only - General

Additional Information:

`use_authtok` is a required `pam_pwhistory(8)` module parameter. It instructs `pam_pwhistory` to accept the password token passed by the preceding module in the PAM stack (typically `pam_unix`) rather than issuing its own authentication prompt. This is distinct from `pwhistory.conf(5)` settings, which govern the history file location and the `remember` count.

Removing `use_authtok` from the `pam_pwhistory.so` line in `/etc/pam.d/common-password` — or omitting it from `/etc/security/pwhistory.conf` — will cause `pam_pwhistory` to prompt for its own password, breaking the password change flow. The `use_authtok` option **must** be present in one of the two supported configuration methods.

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.		●	●

Internal Only - General

5.3.3.4 Configure pam_unix module

The `pam_unix.so` module is the standard Unix authentication module. It uses standard calls from the system's libraries to retrieve and set account information as well as authentication. Usually this is obtained from the `/etc/passwd` and the `/etc/shadow` file as well if shadow is enabled.

Internal Only - General

5.3.3.4.1 Ensure pam_unix does not include nullok (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `nullok` argument overrides the default action of `pam_unix.so` to not permit the user access to a service if their official password is blank.

Rationale:

Using a strong password is essential to helping protect personal and sensitive information from unauthorized access

Audit:

Run the following command to verify that the nullok argument is not set on the pam_unix.so module:

```
# grep -PHs -- '^\h*[^#\n\r]+\h+pam_unix\.so\h+([^#\n\r]+\h+)?nullok\b'
/etc/pam.d/common-{password,auth,account,session,session-noninteractive}
```

Nothing should be returned

Remediation:

Run the following command:

```
# grep -PH -- '^\h*([^#\n\r]+\h+)?pam_unix\.so\h+([^#\n\r]+\h+)?nullok\b'
/usr/share/pam-configs/*
```

Edit any files returned and remove the `nullok` argument for the `pam_unix` lines

Example File:

Internal Only - General

```
Name: Unix authentication
Default: yes
Priority: 256
Auth-Type: Primary
Auth:
        [success=end default=ignore]    pam_unix.so try_first_pass # <-
**ensure line does not include nullok nullok**
Auth-Initial:
        [success=end default=ignore]    pam_unix.so # <- **ensure line does
not include nullok nullok**
Account-Type: Primary
Account:
        [success=end new_authtok_reqd=done default=ignore]      pam_unix.so
Account-Initial:
        [success=end new_authtok_reqd=done default=ignore]      pam_unix.so
Session-Type: Additional
Session:
        required        pam_unix.so
Session-Initial:
        required        pam_unix.so
Password-Type: Primary
Password:
        [success=end default=ignore]    pam_unix.so obscure use_authtok
try_first_pass yescrypt
Password-Initial:
        [success=end default=ignore]    pam_unix.so obscure yescrypt
```

Run the following command to update the files in the `/etc/pam.d/` directory:

```
# pam-auth-update --enable <EDITED_PROFILE_NAME>
```

Example:

```
# pam-auth-update --enable unix
```

Note: If custom files are being used, the corresponding files in `/etc/pam.d/` would need to be edited directly, and the `pam-auth-update --enable <EDITED_PROFILE_NAME>` command skipped

References:

1. GRID: NIX-00000227

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.	●	●	●

Internal Only - General

Controls Version	Control	IG 1	IG 2	IG 3
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.		●	●

Internal Only - General

5.3.3.4.2 Ensure pam_unix does not include remember (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `remember=n` argument saves the last n passwords for each user in `/etc/security/opasswd` in order to force password change history and keep the user from alternating between the same password too frequently. The MD5 password hash algorithm is used for storing the old passwords. Instead of this option the `pam_pwhistory` module should be used. The `pam_pwhistory` module saves the last n passwords for each user in `/etc/security/opasswd` using the password hash algorithm set on the `pam_unix` module. This allows for the `yescrypt` or `sha512` hash algorithm to be used.

Rationale:

The `remember=n` argument should be removed to ensure a strong password hashing algorithm is being used. A stronger hash provides additional protection to the system by increasing the level of effort needed for an attacker to successfully determine local user's old passwords stored in `/etc/security/opasswd`.

Audit:

Run the following command to verify that the `remember` argument is not set on the `pam_unix.so` module:

```
# grep -PHs -- '^\h*[^#\n\r]+\h+pam_unix\.so\h+([^#\n\r]+\h+)?remember=\d+\b'
/etc/pam.d/common-{password,auth,account,session,session-noninteractive}
```

Nothing should be returned

Remediation:

Run the following command:

```
# grep -PH -- '^\h*([^#\n\r]+\h+)?pam_unix\.so\h+([^#\n\r]+\h+)?remember\b'
/usr/share/pam-configs/*
```

Edit any files returned and remove the `remember=_<N>_` argument for the `pam_unix` lines

Example output:

Internal Only - General

```
[success=end default=ignore]   pam_unix.so obscure use_authtok try_first_pass
yescrypt remember=5 # **<- remove remember=<N>**
[success=end default=ignore]   pam_unix.so obscure yescrypt remember=5 # **<-
remove remember=<N>**
```

Run the following command to update the files in the `/etc/pam.d/` directory:

```
# pam-auth-update --enable <EDITED_PROFILE_NAME>
```

Example:

```
# pam-auth-update --enable unix
```

Note: If custom files are being used, the corresponding files in `/etc/pam.d/` would need to be edited directly, and the `pam-auth-update --enable <EDITED_PROFILE_NAME>` command skipped

References:

1. GRID: NIX-00000228

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.		●	●

Internal Only - General

5.3.3.4.3 Ensure pam_unix includes a strong password hashing algorithm (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

A cryptographic hash function converts an arbitrary-length input into a fixed length output. Password hashing performs a one-way transformation of a password, turning the password into another string, called the hashed password.

The `pam_unix` module can be configured to use one of the following hashing algorithms for user's passwords:

- `md5` - When a user changes their password next, encrypt it with the `MD5` algorithm.
- `bigcrypt` - When a user changes their password next, encrypt it with the `DEC C2` algorithm.
- `sha256` - When a user changes their password next, encrypt it with the `SHA256` algorithm. The `SHA256` algorithm must be supported by the crypt(3) function.
- `sha512` - When a user changes their password next, encrypt it with the `SHA512` algorithm. The `SHA512` algorithm must be supported by the crypt(3) function.
- `blowfish` - When a user changes their password next, encrypt it with the `blowfish` algorithm. The `blowfish` algorithm must be supported by the crypt(3) function.
- `gost_yescrypt` - When a user changes their password next, encrypt it with the `gost-yescrypt` algorithm. The `gost-yescrypt` algorithm must be supported by the crypt(3) function.
- `yescrypt` - When a user changes their password next, encrypt it with the `yescrypt` algorithm. The `yescrypt` algorithm must be supported by the crypt(3) function.

Rationale:

The `SHA-512` and `yescrypt` algorithms provide a stronger hash than other algorithms used by Linux for password hash generation. A stronger hash provides additional protection to the system by increasing the level of effort needed for an attacker to successfully determine local user passwords.

Note: These changes only apply to the local system.

Internal Only - General

Audit:

Run the following command to verify that a strong password hashing algorithm is set on the pam_unix.so module:

```
# grep -PH --
'^\h*password\h+([^#\n\r]+)\h+pam_unix\.so\h+([^#\n\r]+\h+)?(sha512|yescrypt)
\b' /etc/pam.d/common-password
```

Output should be similar to:

```
/etc/pam.d/common-password:password   [success=1 default=ignore]
pam_unix.so obscure use_authtok try_first_pass yescrypt
```

Verify that the line(s) include either `sha512` **- OR -** `yescrypt`

Remediation:

Run the following command:

```
# awk '/Password-Type:/{ f = 1;next } /-Type:/{ f = 0 } f {if
(/pam_unix\.so/) print FILENAME}' /usr/share/pam-configs/*
```

Edit any returned files and edit or add a strong hashing algorithm, either `sha512` or `yescrypt`, that meets local site policy to the `pam_unix` lines in the `Password` section:

Example File:

```
Name: Unix authentication
Default: yes
Priority: 256
Auth-Type: Primary # <- Start of "Auth" section
Auth:
        [success=end default=ignore]    pam_unix.so try_first_pass
Auth-Initial:
        [success=end default=ignore]    pam_unix.so
Account-Type: Primary # <- Start of "Account" section
Account:
        [success=end new_authtok_reqd=done default=ignore]      pam_unix.so
Account-Initial:
        [success=end new_authtok_reqd=done default=ignore]      pam_unix.so
Session-Type: Additional  # <- Start of "Session" section
Session:
        required        pam_unix.so
Session-Initial:
        required        pam_unix.so
Password-Type: Primary # <- Start of "Password" section
Password:
        [success=end default=ignore]    pam_unix.so obscure use_authtok
try_first_pass yescrypt # <- **ensure hashing algorithm is either sha512 or
yescrypt**
Password-Initial:
        [success=end default=ignore]    pam_unix.so obscure yescrypt # <-
**ensure hashing algorithm is either sha512 or yescrypt**
```

Run the following command to update the files in the `/etc/pam.d/` directory:

Internal Only - General

```
# pam-auth-update --enable <MODIFIED_PROFILE_NAME>
```

Example:

```
# pam-auth-update --enable unix
```

References:

1. GRID: NIX-00000229
2. NIST SP 800-53 Rev. 5: IA-5

Additional Information:

The following command may be used to expire all non-system user ID's immediately and force them to change their passwords on next login. Any system accounts that need to be expired should be carefully done separately by the system administrator to prevent any potential problems.

```
# awk -F: '( $3<'"$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)"' && $1 !=
"nfsnobody" ) { print $1 }' /etc/passwd | xargs -n 1 chage -d 0
```

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.11 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.		●	●

Internal Only - General

5.3.3.4.4 Ensure pam_unix includes use_authtok (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`use_authtok` - When password changing enforce the module to set the new password to the one provided by a previously stacked password module

Rationale:

`use_authtok` allows multiple pam modules to confirm a new password before it is accepted.

Audit:

Run the following command to verify that `use_authtok` is set on the pam_unix.so module lines in the password stack:

```
# grep -PH --
'^\h*password\h+([^#\n\r]+)\h+pam_unix\.so\h+([^#\n\r]+\h+)?use_authtok\b'
/etc/pam.d/common-password
```

Output should be similar to:

```
/etc/pam.d/common-password:password   [success=1 default=ignore]
pam_unix.so obscure use_authtok try_first_pass yescrypt
```

Verify that the line(s) include `use_authtok`

Remediation:

Run the following command:

```
# awk '/Password-Type:/{ f = 1;next } /-Type:/{ f = 0 } f {if
(/pam_unix\.so/) print FILENAME}' /usr/share/pam-configs/*
```

Edit any returned files add `use_authtok` to the `pam_unix` line in the `Password` section under `Password:` subsection:

Note: The if the file's `Password` section includes a `Password-Initial:` subsection, `use_authtok` should not be added to the `pam_unix` line in the `Password-Initial:` subsection

Example File:

Internal Only - General

```
Name: Unix authentication
Default: yes
Priority: 256
Auth-Type: Primary # <- Start of "Auth" section
Auth:
        [success=end default=ignore]    pam_unix.so try_first_pass
Auth-Initial:
        [success=end default=ignore]    pam_unix.so
Account-Type: Primary # <- Start of "Account" section
Account:
        [success=end new_authtok_reqd=done default=ignore]      pam_unix.so
Account-Initial:
        [success=end new_authtok_reqd=done default=ignore]      pam_unix.so
Session-Type: Additional  # <- Start of "Session" section
Session:
        required        pam_unix.so
Session-Initial:
        required        pam_unix.so
Password-Type: Primary # <- Start of "Password" section
Password:
        [success=end default=ignore]    pam_unix.so obscure use_authtok
try_first_pass yescrypt # <- **ensure line includes use_authtok**
Password-Initial:
        [success=end default=ignore]    pam_unix.so obscure yescrypt # <-
**Password-Initial: subsection does not include use_authtok
```

Run the following command to update the files in the `/etc/pam.d/` directory:

```
# pam-auth-update --enable <MODIFIED_PROFILE_NAME>
```

Example:

```
# pam-auth-update --enable unix
```

References:

1. GRID: NIX-00000230
2. NIST SP 800-53 Rev. 5: IA-5

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.		●	●

Internal Only - General

5.4 User Accounts and Environment

This section provides guidance on setting up secure defaults for system and user accounts and their environment.

Internal Only - General

5.4.1 Configure 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.

Internal Only - General

5.4.1.1 Ensure password expiration is configured (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.

`PASS_MAX_DAYS` <*N*> - The maximum number of days a password may be used. If the password is older than this, a password change will be forced. If not specified, -1 will be assumed (which disables the restriction).

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.

We recommend a yearly password change. This is primarily because for all their good intentions users will share credentials across accounts. Therefore, even if a breach is publicly identified, the user may not see this notification, or forget they have an account on that site. This could leave a shared credential vulnerable indefinitely. Having an organizational policy of a 1-year (annual) password expiration is a reasonable compromise to mitigate this with minimal user burden.

Impact:

The password expiration must be greater than the minimum days between password changes or users will be unable to change their password.

Excessive password expiration requirements do more harm than good, because these requirements make users select predictable passwords, composed of sequential words and numbers that are closely related to each other. In these cases, the next password can be predicted based on the previous one (incrementing a number used in the password for example). Also, password expiration requirements offer no containment benefits because attackers will often use credentials as soon as they compromise them. Instead, immediate password changes should be based on key events including, but not limited to:

- Indication of compromise
- Change of user roles
- When a user leaves the organization.

Internal Only - General

Not only does changing passwords every few weeks or months frustrate the user, but it's also been suggested that it does more harm than good, because it could lead to bad practices by the user such as adding a character to the end of their existing password.

Audit:

Run the following command and verify `PASS_MAX_DAYS` is set to 365 days or less and conforms to local site policy:

```
# grep -Pi -- '^\h*PASS_MAX_DAYS\h+\d+\b' /etc/login.defs
```

Example output:

```
PASS_MAX_DAYS 365
```

The returned value must be between 1 and 365 (inclusive). If `PASS_MAX_DAYS` is absent or the value exceeds 365, remediation is required.

Run the following command to verify all `/etc/shadow` passwords `PASS_MAX_DAYS`:

- is greater than `0` days
- is less than or equal to `365` days
- conforms to local site policy

```
# awk -F: '($2~/^\$.+\$/) {if($5 > 365 || $5 < 1)print "User: " $1 "
PASS_MAX_DAYS: " $5}' /etc/shadow
```

Nothing should be returned

Remediation:

Edit `/etc/login.defs` and set `PASS_MAX_DAYS` to a value greater than `0` that follows local site policy:

Example:

```
PASS_MAX_DAYS 365
```

Run the following command to modify user parameters for all users with a password set to a maximum age no greater than `365` or less than `1` that follows local site policy:

```
# chage --maxdays <N> <user>
```

Example:

```
# awk -F: '($2~/^\$.+\$/) {if($5 > 365 || $5 < 1)system ("chage --maxdays 365
" $1)}' /etc/shadow
```

Warning: If a password has been set at system install or kickstart, the `last change date` field is not set, In this case, setting `PASS_MAX_DAYS` will immediately expire the password. One possible solution is to populate the `last change date` field through a command like: `chage -d "$(date +%Y-%m-%d)" root`

Internal Only - General

Default Value:

PASS_MAX_DAYS 99999

References:

1. GRID: NIX-00000245
2. CIS Password Policy Guide
3. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

Additional Information:

A value of -1 will disable password expiration.

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.		●	●

Internal Only - General

5.4.1.2 Ensure minimum password days is configured (Manual)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

`PASS_MIN_DAYS` <*N*> - The minimum number of days allowed between password changes. Any password changes attempted sooner than this will be rejected. If not specified, 0 will be assumed (which disables the restriction).

Rationale:

Users may have favorite passwords that they like to use because they are easy to remember and they believe that their password choice is secure from compromise. Unfortunately, passwords are compromised and if an attacker is targeting a specific individual user account, with foreknowledge of data about that user, reuse of old, potentially compromised passwords, may cause a security breach.

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

Impact:

If a users password is set by other personnel as a procedure in dealing with a lost or expired password, the user should be forced to update this "set" password with their own password. e.g. force "change at next logon".

If it is not possible to have a user set their own password immediately, and this recommendation or local site procedure may cause a user to continue using a third party generated password, `PASS_MIN_DAYS` for the effected user should be temporally changed to `0`via `chage --mindays <user>`, to allow a user to change their password immediately.

For applications where the user is not using the password at console, the ability to "change at next logon" may be limited. This may cause a user to continue to use a password created by other personnel.

Audit:

Run the following command to verify that `PASS_MIN_DAYS` is set to a value greater than `0` and follows local site policy:

```
# grep -Pi -- '^\h*PASS_MIN_DAYS\h+\d+\b' /etc/login.defs
```

Example output:

Internal Only - General

```
PASS_MIN_DAYS   1
```

Run the following command to verify all passwords have a `PASS_MIN_DAYS` greater than `0`:

```
# awk -F: '($2~/^\$.+\$/) {if($4 < 1)print "User: " $1 " PASS_MIN_DAYS: "
$4}' /etc/shadow
```

Nothing should be returned

Remediation:

Edit `/etc/login.defs` and set `PASS_MIN_DAYS` to a value greater than `0` that follows local site policy:

Example:

```
PASS_MIN_DAYS 1
```

Run the following command to modify user parameters for all users with a password set to a minimum days greater than zero that follows local site policy:

```
# chage --mindays <N> <user>
```

Example:

```
# awk -F: '($2~/^\$.+\$/) {if($4 < 1)system ("chage --mindays 1 " $1)}'
/etc/shadow
```

Default Value:

PASS_MIN_DAYS 0

References:

1. GRID: NIX-00000246
2. CIS Password Policy Guide
3. NIST SP 800-53A :: IA-5 (1).1 (v)
4. NIST SP 800-53 :: IA-5 (1) (d)
5. STIG ID: RHEL-08-020190 | Rule ID: SV-230365r858727 | CAT II
6. STIG ID: UBTU-20-010007 | Rule ID: SV-238202r1015140 | CAT III
7. STIG ID: UBTU-22-411025 | Rule ID: SV-260545r1015007 | CAT III
8. STIG ID: RHEL-08-020180 | Rule ID: SV-230364r627750 | CAT II

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.	●	●	●

Internal Only - General

Controls Version	Control	IG 1	IG 2	IG 3
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.		●	●

Internal Only - General

5.4.1.3 Ensure password expiration warning days is configured (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.

PASS_WARN_AGE <*N*> - The number of days warning given before a password expires. A zero means warning is given only upon the day of expiration, a negative value means no warning is given. If not specified, no warning will be provided.

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.

Impact:

Users warned in advance of password expiration are more likely to set a strong replacement password proactively, rather than being forced to create one immediately under time pressure.

Audit:

Run the following command and verify PASS_WARN_AGE is 7 or more and follows local site policy:

```
# grep -Pi -- '^\h*PASS_WARN_AGE\h+\d+\b' /etc/login.defs
```

Example output:

```
PASS_WARN_AGE 7
```

The returned value must be 7 or more. If PASS_WARN_AGE is absent or the value is less than 7, remediation is required.

Run the following command to verify all passwords have a PASS_WARN_AGE of 7 or more:

```
# awk -F: '($2~/^\$.+\$/) {if($6 < 7)print "User: " $1 " PASS_WARN_AGE: "
$6}' /etc/shadow
```

Nothing should be returned

Internal Only - General

Remediation:

Edit `/etc/login.defs` and set `PASS_WARN_AGE` to a value of `7` or more that follows local site policy:

Example:

```
PASS_WARN_AGE 7
```

Run the following command to modify user parameters for all users with a password set to a minimum warning to `7` or more days that follows local site policy:

```
# chage --warndays <N> <user>
```

Example:

```
# awk -F: '($2~/^\$.+\$/) {if($6 < 7)system ("chage --warndays 7 " $1)}'
/etc/shadow
```

Default Value:

PASS_WARN_AGE 7

References:

1. GRID: NIX-00000247

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.		●	●

Internal Only - General

5.4.1.4 Ensure strong password hashing algorithm is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

A cryptographic hash function converts an arbitrary-length input into a fixed length output. Password hashing performs a one-way transformation of a password, turning the password into another string, called the hashed password.

`ENCRYPT_METHOD` (string) - This defines the system default encryption algorithm for encrypting passwords (if no algorithm are specified on the command line). It can take one of these values:

- `MD5` - MD5-based algorithm will be used for encrypting password
- `SHA256` - SHA256-based algorithm will be used for encrypting password
- `SHA512` - SHA512-based algorithm will be used for encrypting password
- `BCRYPT` - BCRYPT-based algorithm will be used for encrypting password
- `YESCRYPT` - YESCRYPT-based algorithm will be used for encrypting password
- `DES` - DES-based algorithm will be used for encrypting password (default)

Note:

- This parameter overrides the deprecated `MD5_CRYPT_ENAB` variable.
- This parameter will only affect the generation of group passwords.
- The generation of user passwords is done by PAM and subject to the PAM configuration.
- It is recommended to set this variable consistently with the PAM configuration.

Rationale:

The `SHA-512` and `yescrypt` algorithms provide a stronger hash than other algorithms used by Linux for password hash generation. A stronger hash provides additional protection to the system by increasing the level of effort needed for an attacker to successfully determine local group passwords.

Impact:

Changing the hashing algorithm only affects new or changed passwords. Existing passwords continue to use the algorithm in use at the time they were set. A hash algorithm migration requires all users to change their passwords.

Internal Only - General

Audit:

Run the following command to verify the hashing algorithm is `sha512` or `yescrypt` in `/etc/login.defs`:

```
# grep -Pi -- '^\h*ENCRYPT_METHOD\h+(SHA512|YESCRYPT)\b' /etc/login.defs
```

Example output:

```
ENCRYPT_METHOD SHA512
 - OR -
ENCRYPT_METHOD YESCRYPT
```

Remediation:

Edit `/etc/login.defs` and set the `ENCRYPT_METHOD` to `SHA512` or `YESCRYPT`:

```
ENCRYPT_METHOD <HASHING_ALGORITHM>
```

Example:

```
ENCRYPT_METHOD YESCRYPT
```

Note:

- This only effects local groups' passwords created after updating the file to use `sha512` or `yescrypt`.
- If it is determined that the password algorithm being used is not `sha512` or `yescrypt`, once it is changed, it is recommended that all group passwords be updated to use the stronger hashing algorithm.
- It is recommended that the chosen hashing algorithm is consistent across `/etc/login.defs` and the PAM configuration

Default Value:

ENCRYPT_METHOD YESCRYPT

References:

1. GRID: NIX-00000248
2. NIST SP 800-53 Rev. 5: IA-5

Internal Only - General

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.		●	●

Internal Only - General

5.4.1.5 Ensure inactive password lock is configured (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.

`INACTIVE` - Defines the number of days after the password exceeded its maximum age where the user is expected to replace this password.

The value is stored in the shadow password file. An input of `0` will disable an expired password with no delay. An input of `-1` will blank the respective field in the shadow password file.

Rationale:

Inactive accounts pose a threat to system security since the users are not logging in to notice failed login attempts or other anomalies.

Impact:

Locking accounts after extended inactivity limits the attack surface from dormant credentials. Setting `unlock_time=0` (never) for inactive accounts increases the administrative burden when a legitimate user returns after a long absence.

Audit:

Run the following command and verify `INACTIVE` conforms to site policy (no more than 45 days):

```
# useradd -D | grep INACTIVE

INACTIVE=45
```

Run the following command and Review list of users and `INACTIVE` to verify that all users `INACTIVE` conforms to site policy (no more than 45 days):

```
# awk -F: '($2~/^\$.+\$/) {if($7 > 45 || $7 < 0)print "User: " $1 " INACTIVE:
" $7 "Days"}' /etc/shadow
```

Nothing should be returned

Remediation:

Run the following command to set the default password inactivity period to 45 days or less that meets local site policy:

Internal Only - General

```
# useradd -D -f <N>
```

Example:

```
# useradd -D -f 45
```

Run the following command to modify user parameters for all users with a password set to a inactive age of 45 days or less that follows local site policy:

```
# chage --inactive <N> <user>
```

Example:

```
# awk -F: '($2~/^\$.+\$/) {if($7 > 45 || $7 < 0)system ("chage --inactive 45
" $1)}' /etc/shadow
```

Default Value:

INACTIVE=-1

References:

1. GRID: NIX-00000249
2. CIS Password Policy Guide

Additional Information:

A value of -1 would disable this setting.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.2 Use Unique Passwords Use unique passwords for all enterprise assets. Best practice implementation includes, at a minimum, an 8-character password for accounts using MFA and a 14-character password for accounts not using MFA.	●	●	●
v7	4.4 Use Unique Passwords Where multi-factor authentication is not supported (such as local administrator, root, or service accounts), accounts will use passwords that are unique to that system.		●	●

Internal Only - General

5.4.1.6 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 user's recorded password change date is in the future, then they could bypass any set password expiration.

Impact:

An account with a future last-change date can permanently bypass password expiration policies — the shadow system calculates expiry relative to the last-change date, so a far-future date means the password never expires regardless of `PASS_MAX_DAYS`.

Audit:

Run the following script and verify nothing is returned:

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

{
 while IFS= read -r l_user; do
 l_change_raw=$(chage --list "$l_user" | grep '^Last password change' | cut -
d: -f2 | tr -d ' ')
 [ "$l_change_raw" = "never" ] && continue
 l_change=$(date -u -d "$l_change_raw" +%s 2>/dev/null) || continue
 if [[ "$l_change" -gt "$(date +%s)" ]]; then
 echo "User: \"$l_user\" last password change was \"$l_change_raw\""
 fi
 done < <(awk -F: '$2~/^\$.+\$/{print $1}' /etc/shadow)
}
```

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.

Examples

Reset last change date to today for a single user:

Internal Only - General

```
# chage -d "$(date +%Y-%m-%d)" <user>
```

Force password change on next login (recommended companion):

```
# chage -d 0 <user>
```

References:

1. GRID: NIX-00000250

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.		●	●

Internal Only - General

5.4.2 Configure root and system accounts and environment

Internal Only - General

5.4.2.1 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 the Recommendation "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:

Run the following command to change the `root` account UID to `0`:

```
# usermod -u 0 root
```

Modify any users other than `root` with UID `0` and assign them a new UID.

References:

1. GRID: NIX-00000255
2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	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.	●	●	●

Internal Only - General

Controls Version	Control	IG 1	IG 2	IG 3
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.	●	●	●

Internal Only - General

5.4.2.2 Ensure root is the only GID 0 account (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `usermod` command can be used to specify which group the `root` account belongs to. This affects permissions of files that are created by the `root` account.

Rationale:

Using GID 0 for the `root` account helps prevent `root` -owned files from accidentally becoming accessible to non-privileged users.

Audit:

Run the following command to verify the `root` user's primary GID is `0`, and no other user's have GID `0` as their primary GID:

```
# awk -F: '($1 !~ /^(sync|shutdown|halt|operator)/ && $4=="0") {print
$1":"$4}' /etc/passwd

root:0
```

Note: User's: sync, shutdown, halt, and operator are excluded from the check for other user's with GID `0`

Remediation:

Run the following command to set the `root` user's GID to `0`:

```
# usermod -g 0 root
```

Run the following command to set the `root` group's GID to `0`:

```
# groupmod -g 0 root
```

Remove any users other than the `root` user with GID 0 or assign them a new GID if appropriate.

References:

1. GRID: NIX-00000254
2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

Internal Only - General

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.	●	●	●

Internal Only - General

5.4.2.3 Ensure group root is the only GID 0 group (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `groupmod` command can be used to specify which group the `root` group belongs to. This affects permissions of files that are group owned by the `root` group.

Rationale:

Using GID 0 for the `root` group helps prevent `root` group owned files from accidentally becoming accessible to non-privileged users.

Audit:

Run the following command to verify no group other than `root` is assigned GID `0`:

```
# awk -F: '$3=="0"{print $1":"$3}' /etc/group

root:0
```

Remediation:

Run the following command to set the `root` group's GID to `0`:

```
# groupmod -g 0 root
```

Remove any groups other than the `root` group with GID 0 or assign them a new GID if appropriate.

References:

1. GRID: NIX-00000252
2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	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.	●	●	●

Internal Only - General

Controls Version	Control	IG 1	IG 2	IG 3
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.	●	●	●

Internal Only - General

5.4.2.4 Ensure root account access is controlled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

There are a number of methods to access the root account directly. Without a password set any user would be able to gain access and thus control over the entire system.

Rationale:

Access to `root` should be secured at all times.

Impact:

If there are any automated processes that relies on access to the root account without authentication, they will fail after remediation.

Audit:

Run the following command to verify that either the root user's password is set or the root user's account is locked:

```
# passwd -S root | awk '$2 ~ /^(P|L)/ {print "User: \"" $1 "\" Password is
status: " $2}'
```

Verify the output is either:

```
User: "root" Password is status: P
- OR -
User: "root" Password is status: L
```

Note:

- `P` - Password is set
- `L` - Password is locked
- `NP` - No password set

Warning: If no output is returned, the root account is configured with no password (`NP` status) — this is non-compliant and remediation is required.

Remediation:

Run the following command to set a password for the `root` user:

```
# passwd root
```

- OR -

Internal Only - General

Run the following command to lock the `root` user account:

```
# usermod -L root
```

References:

1. GRID: NIX-00000253

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.	●	●	●

Internal Only - General

5.4.2.5 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.

Impact:

A compromised or misconfigured root PATH can cause the root user to inadvertently execute a trojan program placed in a world-writable directory — this is a classic privilege escalation vector.

Audit:

Run the following script to verify root's path does not include:

- Locations that are not directories
- An empty directory (`::`)
- A trailing (`:`)
- Current working directory (`.`)
- Non `root` owned directories
- Directories that less restrictive than mode `0755`

Internal Only - General

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

{
 a_output2=() l_pmask="0022"
 l_maxperm="$( printf '%o' $(( 0777 & ~$l_pmask )) )"
 l_root_path="$(sudo su - root -c env | awk -F= '$1=="PATH"{print $2}')"
 IFS=":" read -ra a_path_loc <<< "$l_root_path"
 grep -q -- "::" <<< "$l_root_path" && \
 a_output2+=(" - root's path contains a empty directory (::)")
 grep -Pq -- ":\h*$" <<< "$l_root_path" && \
 a_output2+=(" - root's path contains a trailing (:)")
 grep -Pq -- '(^\h*|:)\.(:|\h*$)' <<< "$l_root_path" && \
 a_output2+=(" - root's path contains current working directory (.)")
 for l_path in "${a_path_loc[@]}"; do
 if [ -d "$l_path" ]; then
 while IFS=: read -r l_fmode l_fown; do
 [ "$l_fown" != "root" ] && \
 a_output2+=(" - Directory: \"$l_path\" is owned by: \"$l_fown\"" \
 " should be owned by \"root\"")
 [ $(( $l_fmode & $l_pmask )) -gt 0 ] && \
 a_output2+=(" - Directory: \"$l_path\" is mode: \"$l_fmode\"" \
 " and should be mode: \"$l_maxperm\" or more restrictive")
 done <<< "$(stat -Lc '%#a:%U' "$l_path")"
 else
 a_output2+=(" - \"$l_path\" is not a directory")
 fi
 done
 if [ "${#a_output2[@]}" -gt 0 ]; then
 printf '%s\n' "Non-compliant path elements:" "${a_output2[@]}"
 fi
}
```

Nothing should be returned. Any output indicates a non-compliant path element in root's PATH.

Remediation:

Correct or justify any:

- Locations that are not directories
- Empty directories (`::`)
- Trailing (`:`)
- Current working directory (`.`)
- Non `root` owned directories
- Directories that less restrictive than mode `0755`

References:

1. GRID: NIX-00000256
2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

5.4.2.6 Ensure root user umask is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The user file-creation mode mask (`umask`) is used to determine the file permission for newly created directories and files. In Linux, the default permissions for any newly created directory is 0777 (`rwxrwxrwx`), and for any newly created file it is 0666 (`rw-rw-rw-`). The `umask` modifies the default Linux permissions by restricting (masking) these permissions. The `umask` is not simply subtracted, but is processed bitwise. Bits set in the `umask` are cleared in the resulting file mode.

`umask` can be set with either `Octal` or `Symbolic` values:

- `Octal` (Numeric) Value - Represented by either three or four digits. ie `umask 0027` or `umask 027`. If a four digit umask is used, the first digit is ignored. The remaining three digits effect the resulting permissions for user, group, and world/other respectively.
- `Symbolic` Value - Represented by a comma separated list for User `u`, group `g`, and world/other `o`. The permissions listed are not masked by `umask`. ie a `umask` set by `umask u=rwx,g=rx,o=` is the `Symbolic` equivalent of the `Octal umask 027`. This `umask` would set a newly created directory with file mode `drwxr-x---` and a newly created file with file mode `rw-r-----`.

root user Shell Configuration Files:

- `/root/.profile` - Is executed to configure the root users' shell before the initial command prompt. **Is only read by login shells.**
- `/root/.bashrc` - Is executed for interactive shells. **only read by a shell that's both interactive and non-login**

`umask` is set by order of precedence. If `umask` is set in multiple locations, this order of precedence will determine the system's default `umask`.

Order of precedence:

1. `/root/.profile`
2. `/root/.bashrc`
3. The system default umask

Internal Only - General

Rationale:

Setting a secure value for `umask` ensures that users make a conscious choice about their file permissions. A permissive `umask` value could result in directories or files with excessive permissions that can be read and/or written to by unauthorized users.

Audit:

Run the following command to verify that the `root` user files, `/root/.profile` and `/root/.bashrc`, do not include a `umask` that is less restrictive than `0027`:

```
# grep -Psi -- '^\h*umask\h+((\d{1,2}(\d[^7]|[^2-
7]\d)\b)|(u=[rwx]{1,3},)?(((g=[rx]?[rx]?w[rx]?[rx]?\b)(,o=[rwx]{1,3})?)|((g=[
wrx]{1,3},)?o=[wrx]{1,3}\b)))' /root/.profile /root/.bashrc
```

Nothing should be returned.

Remediation:

Edit `/root/.profile` and `/root/.bashrc` and either:

- remove, comment out, or update any line with `umask`.

- OR -

- update any line that includes `umask` to a value of `0027` or more restrictive.

Example:

```
umask 027
```

Note: the Recommendation "Ensure default user umask is configured" includes guidance to set the default umask

Default Value:

System default `umask`

References:

1. GRID: NIX-00000257
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

Internal Only - General

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.	●	●	●

Internal Only - General

5.4.2.7 Ensure system accounts do not have a valid login shell (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. Furthermore, a user may add special accounts that are not intended to provide an interactive shell.

Rationale:

It is important to make sure that accounts that are not being used by regular users are prevented from being used to provide an interactive shell. By default, most distributions set the password field for these accounts to an invalid string, but it is also recommended that the shell field in the password file be set to the `nologin` shell. This prevents the account from potentially being used to run any commands.

Audit:

Run the following command to verify system accounts, except for `root`, `halt`, `sync`, `shutdown` or `nfsnobody`, do not have a valid login shell:

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

{
   l_valid_shells="^($(awk -F\/ '$NF != "nologin" {print}' /etc/shells | sed
-rn '/^\//{s,/,\\\\/,g;p}' | paste -s -d '|' - ))$"
   awk -v pat="$l_valid_shells" -F:
'($1!~/^(root|halt|sync|shutdown|nfsnobody)$/ && ($3<'"$(awk
'/^\s*UID_MIN/{print $2}' /etc/login.defs)"' || $3 == 65534) && $(NF) ~ pat)
{print "Service account: \"" $1 "\" has a valid shell: " $7}' /etc/passwd
}
```

Nothing should be returned

Remediation:

Run the following command to set the shell for any service accounts returned by the audit to `nologin`:

```
# usermod -s $(command -v nologin) <user>
```

Example script:

Internal Only - General

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

{
   l_valid_shells="^($( awk -F\/ '$NF != "nologin" {print}' /etc/shells | sed
-rn '/^\//{s,/,\\\\/,g;p}' | paste -s -d '|' - ))$"
   awk -v pat="$l_valid_shells" -F:
'($1!~/^(root|halt|sync|shutdown|nfsnobody)$/ && ($3<'"$(awk
'/^\s*UID_MIN/{print $2}' /etc/login.defs)"' || $3 == 65534) && $(NF) ~ pat)
{system ("usermod -s '"$(command -v nologin)"' " $1)}' /etc/passwd
}
```

References:

1. GRID: NIX-00000258
2. NIST SP 800-53 Rev. 5: AC-2(5), AC-3, AC-11, MP-2

Additional Information:

The `root`, `sync`, `shutdown`, and `halt` users are exempted from requiring a non-login shell.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 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.	●	●	●

Internal Only - General

5.4.2.8 Ensure accounts without a valid login shell are locked (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. Furthermore, a user may add special accounts that are not intended to provide an interactive shell.

Rationale:

It is important to make sure that accounts that are not being used by regular users are prevented from being used to provide an interactive shell. By default, most distributions set the password field for these accounts to an invalid string, but it is also recommended that the shell field in the password file be set to the `nologin` shell. This prevents the account from potentially being used to run any commands.

Audit:

Run the following script to verify all non-root accounts without a valid login shell are locked.

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

{
   l_valid_shells="^($(awk -F\/ '$NF != "nologin" {print}' /etc/shells | sed
-rn '/^\//{s,/,\\\\/,g;p}' | paste -s -d '|' - ))$"
   while IFS= read -r l_user; do
      passwd -S "$l_user" | awk '$2 !~ /^L/ {print "Account: \"" $1 "\" does
not have a valid login shell and is not locked"}'
   done < <(awk -v pat="$l_valid_shells" -F: '($1 != "root" && $(NF) !~ pat)
{print $1}' /etc/passwd)
}
```

Nothing should be returned

Remediation:

Run the following command to lock any non-root accounts without a valid login shell returned by the audit:

```
# usermod -L <user>
```

Example script::

Internal Only - General

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

{
   l_valid_shells="^($(awk -F\/ '$NF != "nologin" {print}' /etc/shells | sed
-rn '/^\//{s,/,\\\\/,g;p}' | paste -s -d '|' - ))$"
   while IFS= read -r l_user; do
      passwd -S "$l_user" | awk '$2 !~ /^L/ {system ("usermod -L " $1)}'
   done < <(awk -v pat="$l_valid_shells" -F: '($1 != "root" && $(NF) !~ pat)
{print $1}' /etc/passwd)
}
```

References:

1. GRID: NIX-00000251
2. NIST SP 800-53 Rev. 5: AC-2(5), AC-3, AC-11, MP-2

CIS Controls:

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

Internal Only - General

5.4.3 Configure user default environment

Internal Only - General

5.4.3.1 Ensure nologin is not listed in /etc/shells (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

`/etc/shells` is a text file which contains the full pathnames of valid login shells. This file is consulted by `chsh` and available to be queried by other programs.

Be aware that there are programs which consult this file to find out if a user is a normal user; for example, FTP daemons traditionally disallow access to users with shells not included in this file.

Rationale:

A user can use `chsh` to change their configured shell.

If a user has a shell configured that isn't in in `/etc/shells`, then the system assumes that they're somehow restricted. In the case of `chsh` it means that the user cannot change that value.

Other programs might query that list and apply similar restrictions.

By putting `nologin` in `/etc/shells`, any user that has `nologin` as its shell is considered a full, unrestricted user. This is not the expected behavior for `nologin`.

Audit:

Run the following command to verify that `nologin` is not listed in the `/etc/shells` file:

```
# grep -Ps '^\h*([^#\n\r]+)?\/nologin\b' /etc/shells
```

Nothing should be returned

Remediation:

Edit `/etc/shells` and remove any lines that include `nologin`

References:

1. GRID: NIX-00000259
2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5
3. shells(5)

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

5.4.3.2 Ensure default user shell timeout is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`TMOUT` is an environmental setting that determines the timeout of a shell in seconds.

- TMOUT=*n* - Sets the shell timeout to *n* seconds. A setting of `TMOUT=0` disables timeout.
- readonly TMOUT- Sets the TMOUT environmental variable as readonly, preventing unwanted modification during run-time.
- export TMOUT - exports the TMOUT variable

System Wide Shell Configuration Files:

- `/etc/profile` - used to set system wide environmental variables on users shells. The variables are sometimes the same ones that are in the `.bash_profile`, however this file is used to set an initial PATH or PS1 for all shell users of the system. **is only executed for interactive *login* shells, or shells executed with the --login parameter.**
- `/etc/profile.d` - `/etc/profile` will execute the scripts within `/etc/profile.d/*.sh`. It is recommended to place your configuration in a shell script within `/etc/profile.d` to set your own system wide environmental variables.
- `/etc/bash.bashrc` - System wide version of `.bashrc`. In Fedora derived distributions, `/etc/bashrc` also invokes /etc/profile.d/*.sh if *non-login* shell, but redirects output to `/dev/null` if *non-interactive.* **Is only executed for *interactive* shells or if `BASH_ENV` is set to `/etc/bash.bashrc`.**

Rationale:

Setting a timeout value reduces the window of opportunity for unauthorized user access to another user's shell session that has been left unattended. It also ends the inactive session and releases the resources associated with that session.

Impact:

Configuring `TMOUT` as `readonly` and `exported` ensures it cannot be overridden by user-defined shell functions or per-user `.bashrc` settings. If `TMOUT` is set without `readonly`, users can trivially disable it with `TMOUT=0`.

Internal Only - General

Audit:

Run the following script to verify that `TMOUT` is configured to: include a timeout of no more than `900` seconds, to be `readonly`, to be `exported`, and is not being changed to a longer timeout.

Internal Only - General

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

{
 a_output=(); a_output2=(); l_tmout_set="900"
 f_tmout_read_chk()
 {
 a_out=(); a_out2=()
 l_tmout_readonly="$(grep -P -- '^\h*(typeset\h\-
xr\hTMOUT=\d+|([^#\n\r]+)?\breadonly\h+TMOUT\b)' "$l_file")"
 l_tmout_export="$(grep -P -- '^\h*(typeset\h\-
xr\hTMOUT=\d+|([^#\n\r]+)?\bexport\b([^#\n\r]+\b)?TMOUT\b)' "$l_file")"
 if [ -n "$l_tmout_readonly" ]; then
 a_out+=(" - Readonly is set as: \"$l_tmout_readonly\" in: \"$l_file\"")
 else
 a_out2+=(" - Readonly is not set in: \"$l_file\"")
 fi
 if [ -n "$l_tmout_export" ]; then
 a_out+=(" - Export is set as: \"$l_tmout_export\" in: \"$l_file\"")
 else
 a_out2+=(" - Export is not set in: \"$l_file\"")
 fi
 }
 while IFS= read -r l_file; do
 l_tmout_value="$(grep -Po -- '^([^#\n\r]+)?\bTMOUT=\d+\b' "$l_file" | awk -
F= '{print $2}')"
 f_tmout_read_chk
 if [ -n "$l_tmout_value" ]; then
 if [[ "$l_tmout_value" -le "$l_tmout_set" && "$l_tmout_value" -gt "0" ]];
then
 a_output+=(" - TMOUT is set to: \"$l_tmout_value\" in: \"$l_file\"")
 [ "${#a_out[@]}" -gt 0 ] && a_output+=("${a_out[@]}")
 [ "${#a_out2[@]}" -gt 0 ] && a_output2+=("${a_out[@]}")
 fi
 if [[ "$l_tmout_value" -gt "$l_tmout_set" || "$l_tmout_value" -le "0" ]];
then
 a_output2+=(" - TMOUT is incorrectly set to: \"$l_tmout_value\" in:
\"$l_file\"")
 [ "${#a_out[@]}" -gt 0 ] && a_output2+=(" ** Incorrect TMOUT value **"
"${a_out[@]}")
 [ "${#a_out2[@]}" -gt 0 ] && a_output2+=("${a_out2[@]}")
 fi
 else
 [ "${#a_out[@]}" -gt 0 ] && a_output2+=(" - TMOUT is not set" "${a_out[@]}")
 [ "${#a_out2[@]}" -gt 0 ] && a_output2+=(" - TMOUT is not set"
"${a_out2[@]}")
 fi
 done < <(grep -Pls -- '^([^#\n\r]+)?\bTMOUT\b' /etc/*bashrc /etc/profile
/etc/profile.d/*.sh)
 [[ "${#a_output[@]}" -le 0 && "${#a_output2[@]}" -le 0 ]] && a_output2+=(" -
TMOUT is not configured")
 if [ "${#a_output2[@]}" -gt 0 ]; then
 printf '%s\n' "Non-compliant TMOUT configuration:" "${a_output2[@]}"
 fi
}
```

Nothing should be returned. Any output identifies a file where `TMOUT` is missing, incorrectly set, not marked readonly, or not exported.

Internal Only - General

Note: If `TMOUT` is set as `readonly` through `readonly TMOUT` and/or `typeset -xr` in more than once, you will receive an error message when logging into a terminal session or connecting with openSSH. It is recommended that `TMOUT` be set only once in **only one** file.

Remediation:

Review `/etc/bashrc`, `/etc/profile`, and all files ending in `*.sh` in the `/etc/profile.d/` directory and remove or edit all `TMOUT=_n_` entries to follow local site policy. `TMOUT` should not exceed 900 or be equal to `0`.

Configure `TMOUT` in **one** of the following files:

- A file in the `/etc/profile.d/` directory ending in `.sh`
- `/etc/profile`
- `/etc/bashrc`

Example command to set TMOUT to `900` seconds in a file in `/etc/profile.d/`:

```
# printf '%s\n' "# Set TMOUT to 900 seconds" "typeset -xr TMOUT=900" >
/etc/profile.d/50-tmout.sh
```

`TMOUT` configuration examples:

```
typeset -xr TMOUT=900
```

Deprecated methods:

- As multiple lines:

```
TMOUT=900
readonly TMOUT
export TMOUT
```

- As a single line:

```
readonly TMOUT=900 ; export TMOUT
```

References:

1. GRID: NIX-00000260

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.

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.3 Configure Automatic Session Locking on Enterprise Assets Configure automatic session locking on enterprise assets after a defined period of inactivity. For general purpose operating systems, the period must not exceed 15 minutes. For mobile end-user devices, the period must not exceed 2 minutes.	●	●	●
v7	16.11 Lock Workstation Sessions After Inactivity Automatically lock workstation sessions after a standard period of inactivity.	●	●	●

Internal Only - General

5.4.3.3 Ensure default user umask is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The user file-creation mode mask (`umask`) is used to determine the file permission for newly created directories and files. In Linux, the default permissions for any newly created directory is 0777 (`rwxrwxrwx`), and for any newly created file it is 0666 (`rw-rw-rw-`). The `umask` modifies the default Linux permissions by restricting (masking) these permissions. The `umask` is not simply subtracted, but is processed bitwise. Bits set in the `umask` are cleared in the resulting file mode.

`umask` can be set with either `Octal` or `Symbolic` values:

- `Octal` (Numeric) Value - Represented by either three or four digits. ie `umask 0027` or `umask 027`. If a four digit umask is used, the first digit is ignored. The remaining three digits effect the resulting permissions for user, group, and world/other respectively.
- `Symbolic` Value - Represented by a comma separated list for User `u`, group `g`, and world/other `o`. The permissions listed are not masked by `umask`. ie a `umask` set by `umask u=rwx,g=rx,o=` is the `Symbolic` equivalent of the `Octal umask 027`. This `umask` would set a newly created directory with file mode `drwxr-x---` and a newly created file with file mode `rw-r-----`.

The default `umask` can be set to use the `pam_umask` module or in a `System Wide Shell Configuration File`. The user creating the directories or files has the discretion of changing the permissions via the chmod command, or choosing a different default `umask` by adding the `umask` command into a `User Shell Configuration File`, (`.bash_profile` or `.bashrc`), in their home directory.

Setting the default umask:

- pam_umask module:
  - will set the umask according to the system default in `/etc/login.defs` and user settings, solving the problem of different `umask` settings with different shells, display managers, remote sessions etc.
  - `umask=<mask>` value in the `/etc/login.defs` file is interpreted as Octal
  - Setting `USERGROUPS_ENAB` to yes in `/etc/login.defs` (default):
    - will enable setting of the `umask` group bits to be the same as owner bits. (examples: 022 -> 002, 077 -> 007) for non-root users, if the `uid` is the same as `gid`, and `username` is the same as the `<primary group name>`

Internal Only - General

        - userdel will remove the user's group if it contains no more members, and useradd will create by default a group with the name of the user
- `System Wide Shell Configuration File`:
    - `/etc/profile` - used to set system wide environmental variables on users shells. The variables are sometimes the same ones that are in the `.bash_profile`, however this file is used to set an initial PATH or PS1 for all shell users of the system. **is only executed for interactive *login* shells, or shells executed with the --login parameter.**
    - `/etc/profile.d` - `/etc/profile` will execute the scripts within `/etc/profile.d/*.sh`. It is recommended to place your configuration in a shell script within `/etc/profile.d` to set your own system wide environmental variables.
    - `/etc/bashrc` - System wide version of `.bashrc`. In Fedora derived distributions, `etc/bashrc` also invokes /etc/profile.d/*.sh if *non-login* shell, but redirects output to `/dev/null` if *non-interactive.* **Is only executed for *interactive* shells or if `BASH_ENV` is set to `/etc/bashrc`.**

User Shell Configuration Files:

- `~/.bash_profile` - Is executed to configure your shell before the initial command prompt. **Is only read by login shells.**
- `~/.bashrc` - Is executed for interactive shells. **only read by a shell that's both interactive and non-login**

`umask` is set by order of precedence. If `umask` is set in multiple locations, this order of precedence will determine the system's default `umask`.

Order of precedence:

1. A file in `/etc/profile.d/` ending in `.sh` - This will override any other system-wide `umask` setting
2. In the file `/etc/profile`
3. On the `pam_umask.so` module in `/etc/pam.d/postlogin`
4. In the file `/etc/login.defs`
5. In the file `/etc/default/login`

Rationale:

Setting a secure default value for `umask` ensures that users make a conscious choice about their file permissions. A permissive `umask` value could result in directories or files with excessive permissions that can be read and/or written to by unauthorized users.

Satisfies: SRG-OS-000480-GPOS-00228, SRG-OS-000480-GPOS-00227

Internal Only - General

Audit:

Run the following commands to verify the systemwide `umask` is `0027` `(u=rwx,g=rx,o=)` or more restrictive:

1. Run the following command to verify `umask` is set to `027` or more restrictive in a file ending in `.sh` in `/etc/profile.d/` and `/etc/profile.d/`

```
# grep -PHsi -- '^\h*umask\h+(?!0?[02367]7\b)' /etc/profile /etc/profile.d/
```

Nothing should be returned.

2. Run the following command to verify `UMASK` is `027` or more restrictive in `/etc/login.defs`:

```
# grep -PHsi -- '^\h*UMASK\h+(?!0?[02367]7\b)' /etc/login.defs
```

Nothing should be returned.

Remediation:

1. Run the following script to comment out all occurrences of `umask` that are less restrictive than `027` in files ending in `*.sh` in the `/etc/profile.d/` directory:

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

{
   while IFS= read -r -d $'\0' l_file; do
      sed -ri '/^\s*umask\s+0?(0[01][0-7]|0[0-7][^7]|[^0][0-7][0-
7])(\s*|\s+.*)$/s/^/# /' "$l_file"
   done < <(find /etc/profile.d/ -type f -name '*.sh' -print0)
}
```

2. Create or edit a file in `/etc/profile.d/` ending in `*.sh` and add or modify the following line:

```
umask 0027
```

Example:

```
# printf '%s\n' "" "umask 027" >> /etc/profile.d/60-default_umask.sh
```

3. Edit `/etc/login.defs` and add or update the following line:

```
UMASK 027
```

Notes:

- This method only applies to bash and shell. If other shells are supported on the system, it is recommended that their configuration files also are checked

Internal Only - General

- If the `pam_umask.so` module is going to be used to set `umask`, ensure that it's not being overridden by another setting. Refer to the PAM_UMASK(8) man page for more information

Default Value:

UMASK 022

References:

1. GRID: NIX-00000261
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-08-020351 | RULE ID: SV-230383r1017192 | CAT II
5. STIG ID: RHEL-08-020353 | RULE ID: SV-230385r1017194 | CAT II
6. STIG ID: RHEL-09-412055 | RULE ID: SV-258072r1045155 | CAT II
7. STIG ID: RHEL-09-412065 | RULE ID: SV-258074r991590 | CAT II
8. STIG ID: RHEL-09-412070 | RULE ID: SV-258075r991590 | CAT II
9. STIG ID: ALMA-09-017620 | RULE ID: SV-269235r1050117 | CAT II
10. STIG ID: ALMA-09-017840 | RULE ID: SV-269237r1050119 | CAT II
11. STIG ID: UBTU-22-412035 | RULE ID: SV-260555r991590 | CAT II
12. STIG ID: UBTU-24-300030 | RULE ID: SV-270716r1066637 | CAT II
13. STIG ID: SLES-15-040420 | RULE ID: SV-235030r991590 | CAT II
14. CCI-000366

Additional Information:

- Other methods of setting a default user umask exist
- If other methods are in use in your environment they should be audited
- The default user umask can be overridden with a user specific umask
- The user creating the directories or files has the discretion of changing the permissions:
    - Using the chmod command
    - Setting a different default umask by adding the umask command into a User Shell Configuration File, (.bashrc), in their home directory
    - Manually changing the umask for the duration of a login session by running the umask command

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 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.	●	●	●

Internal Only - General

Controls Version	Control	IG 1	IG 2	IG 3
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.	●	●	●

6 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.

Internal Only - General

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`).

6.1 System Logging

Logging services should be configured to prevent information leaks and to aggregate logs on a remote server so that they can be reviewed in the event of a system compromise. A centralized log server provides a single point of entry for further analysis, monitoring and filtering.

Security principals for logging

- Ensure transport layer security is implemented between the client and the log server.
- Ensure that logs are rotated as per the environment requirements.
- Ensure all locally generated logs have the appropriate permissions.
- Ensure all security logs are sent to a remote log server.
- Ensure the required events are logged.

What is covered

This section will cover the minimum best practices for the usage of ***either*** `rsyslog` ***or*** `journald`. The recommendations are written such that each is wholly independent of each other and ***only one is implemented***.

- If your organization makes use of an enterprise wide logging system completely outside of `rsyslog` or `journald`, then the following recommendations do not directly apply. However, the principals of the recommendations should be followed regardless of what solution is implemented. If the enterprise solution incorporates either of these tools, careful consideration should be given to the following recommendations to determine exactly what applies.
- Should your organization make use of both `rsyslog` and `journald`, take care how the recommendations may or may not apply to you.

What is not covered

- Enterprise logging systems not utilizing `rsyslog` or `journald`.

Internal Only - General

As logging is very situational and dependent on the local environment, not everything can be covered here.

- Transport layer security should be applied to all remote logging functionality. Both `rsyslog` and `journald` supports secure transport and should be configured as such.
- The log server. There are a multitude of reasons for a centralized log server (and keeping a short period of logging on the local system), but the log server is out of scope for these recommendations.

6.1.1 Configure journald

Included in the systemd suite is a journaling service called `systemd-journald.service` for the collection and storage of logging data. It creates and maintains structured, indexed journals based on logging information that is received from a variety of sources such as:

- Classic RFC3164 BSD syslog via the /dev/log socket
- STDOUT/STDERR of programs via StandardOutput=journal + StandardError=journal in service files (both of which are default settings)
- Kernel log messages via the /dev/kmsg device node
- Audit records via the kernel's audit subsystem
- Structured log messages via journald's native protocol

Any changes made to the `systemd-journald` configuration will require a re-start of `systemd-journald`

Internal Only - General

6.1.1.1 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
- Simple system log messages, via the libc syslog call
- Structured system log messages via the native Journal API
- Standard output and standard error of service units
- Audit records, originating from the kernel audit subsystem

The daemon will implicitly collect numerous metadata fields for each log messages in a secure and unfakeable way. See systemd.journal-fields man page for more information about the collected metadata.

The journal service stores log data either persistently below `/var/log/journal` or in a volatile way below `/run/log/journal/`. By default, log data is stored persistently if `/var/log/journal/` exists during boot, with an implicit fallback to volatile storage. Use `Storage=` in `journald.conf` to configure where log data is placed, independently of the existence of `/var/log/journal/`.

On systems where `/var/log/journal/` does not exist but where persistent logging is desired, and the default `journald.conf` is used, it is sufficient to create the directory and ensure it has the correct access modes and ownership.

Internal Only - General

6.1.1.1.1 Ensure journald service is active (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Ensure that the `systemd-journald` service is enabled to allow capturing of logging events.

Rationale:

If the `systemd-journald` service is not enabled to start on boot, the system will not capture logging events.

Impact:

Restarting systemd-journald may briefly interrupt log collection.

Audit:

Run the following command to verify `systemd-journald` is active:

```
# systemctl is-active systemd-journald.service

active
```

Remediation:

Run the following commands to unmask, enable, and start `systemd-journald.service`

```
# systemctl unmask systemd-journald.service
# systemctl --now enable systemd-journald.service
```

References:

1. GRID: NIX-00000268
2. CCI-001665
3. NIST SP 800-53 :: SC-24
4. NIST SP 800-53A :: SC-24.1 (v)
5. STIG ID: RHEL-09-211040 | RULE ID: SV-257783r991562 | CAT II
6. STIG ID: ALMA-09-041490 | RULE ID: SV-269428r1050311 | CAT II
7. STIG ID: OL09-00-002400 | RULE ID: SV-271739r1091929 | CAT II

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 Collect Audit Logs Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.	●	●	●
v7	6.2 Activate audit logging Ensure that local logging has been enabled on all systems and networking devices.	●	●	●
v7	6.3 Enable Detailed Logging Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.		●	●

Internal Only - General

6.1.1.1.2 Ensure systemd-journal-remote service is not in use (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Journald `systemd-journal-remote` supports the ability to receive messages from remote hosts, thus acting as a log server. Clients should not receive data from other hosts.

Note:

- The same package, `systemd-journal-remote`, is used for both sending logs to remote hosts and receiving incoming logs.
- With regards to receiving logs, there are two services; `systemd-journal-remote.socket` and `systemd-journal-remote.service`.

Rationale:

If a client is configured to also receive data, thus turning it into a server, the client system is acting outside it's operational boundary.

Note: This recommendation **only applies if `journald` is the chosen method for client side logging**. Do not apply this recommendation if `rsyslog` is used.

Impact:

Disabling systemd-journal-remote may break log forwarding if the service was intentionally configured; verify with the operations team before disabling.

Audit:

Run the following command to verify `systemd-journal-remote.socket` and `systemd-journal-remote.service` are not enabled:

```
# systemctl is-enabled systemd-journal-remote.socket systemd-journal-
remote.service | grep -P -- '^enabled'
```

Nothing should be returned

Run the following command to verify `systemd-journal-remote.socket` and `systemd-journal-remote.service` are not active:

```
# systemctl is-active systemd-journal-remote.socket systemd-journal-
remote.service | grep -P -- '^active'
```

Nothing should be returned

Internal Only - General

Remediation:

Run the following commands to stop and mask `systemd-journal-remote.socket` and systemd-journal-remote.service:

```
# systemctl stop systemd-journal-remote.socket systemd-journal-remote.service
# systemctl mask systemd-journal-remote.socket systemd-journal-remote.service
```

References:

1. GRID: NIX-00000283
2. NIST SP 800-53 Rev. 5: AU-2, AU-7 AU-12

CIS Controls:

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

Internal Only - General

6.1.1.1.3 Ensure journald is configured to send logs to rsyslog (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Data from `systemd-journald` may be stored in volatile memory or persisted locally on the server. Utilities exist to accept remote export of `systemd-journald` logs, however, use of the `rsyslog` service provides a consistent means of log collection and export.

Rationale:

- IF - `rsyslog` is the preferred method for capturing logs, all logs of the system should be sent to it for further processing.

> **Note:** This recommendation **only applies if `rsyslog` is the chosen method for client side logging**. Do not apply this recommendation if `systemd-journald` is used.

Audit:

- IF - `rsyslog` is the preferred method for capturing logs

Run the following command to verify journald is configured to forward logs to `rsyslog` (the value should be `yes`):

```
# systemd-analyze cat-config systemd/journald.conf | grep -Pi --
'^\h*ForwardToSyslog\h*=\h*yes\b'
```

`ForwardToSyslog=yes` should be returned.

Run the following command to verify `systemd-journald.service` and `rsyslog.service` are loaded and active:

```
# systemctl list-units --type service | grep -P -- '(journald|rsyslog)'
```

Output should be similar to:

```
rsyslog.service loaded active running System Logging Service
systemd-journald.service loaded active running Journal Service
```

Remediation:

- IF - `Journald` is the preferred method for capturing logs, this section and Recommendation should be skipped and the "Configure Journald" section followed.

- IF - `rsyslog` is the preferred method for capturing logs:

Internal Only - General

Set the following parameter in the `[Journal]` section in `/etc/systemd/journald.conf` or a file in `/etc/systemd/journald.conf.d/` ending in `.conf`:

```
ForwardToSyslog=yes
```

Example:

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

{
 [ ! -d /etc/systemd/journald.conf.d/ ] && mkdir -p
/etc/systemd/journald.conf.d/
 if grep -Psq -- '^\h*\[Journal\]' /etc/systemd/journald.conf.d/60-
journald.conf; then
 printf '%s\n' "ForwardToSyslog=yes" >> /etc/systemd/journald.conf.d/60-
journald.conf
 else
 printf '%s\n' "[Journal]" "ForwardToSyslog=yes" >>
/etc/systemd/journald.conf.d/60-journald.conf
 fi
}
```

Note: Drop-in configuration files have higher precedence and override the main configuration file. Files in the *.conf.d/ configuration subdirectories are sorted by their filename in lexicographic order, regardless of in which of the subdirectories they reside. When multiple files specify the same option, for options which accept just a single value, the entry in the file sorted last takes precedence, and for options which accept a list of values, entries are collected as they occur in the sorted files.

Run to following command to update the parameters in the service:

Restart `systemd-journald.service`:

```
# systemctl reload-or-restart systemd-journald.service
```

References:

1. GRID: NIX-00000265
2. NIST SP 800-53 Rev. 5: AC-3, AU-2, AU-4, AU-12, MP-2
3. SYSTEMD-JOURNALD.SERVICE(8)
4. JOURNALD.CONF(5)

Internal Only - General

Additional Information:

As noted in the `systemd-journald` man pages, `systemd-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.

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.	●	●	●
v8	8.9 Centralize Audit Logs Centralize, to the extent possible, audit log collection and retention across enterprise assets.		●	●
v7	6.2 Activate audit logging Ensure that local logging has been enabled on all systems and networking devices.	●	●	●
v7	6.3 Enable Detailed Logging Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.		●	●
v7	6.5 Central Log Management Ensure that appropriate logs are being aggregated to a central log management system for analysis and review.		●	●

Internal Only - General

6.1.1.1.4 Ensure journald log file access is configured (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Journald will create logfiles that do not already exist on the system. This setting controls what permissions will be applied to these newly created files.

Rationale:

It is important to ensure that log files have the correct permissions to ensure that sensitive data is archived and protected.

Impact:

Restricting journal directory permissions may break tooling that reads log files as a non-root user; audit dependent tooling before applying.

Audit:

First determine if there is an override file `/etc/tmpfiles.d/systemd.conf`. If so, this file overrides all default settings defined in `/usr/lib/tmpfiles.d/systemd.conf` and should be inspected.

If no override file exists, inspect the default `/usr/lib/tmpfiles.d/systemd.conf` against the site specific requirements.

Ensure that the mode of any journald log file referenced by these `tmpfiles.d` entries is `0640` or more restrictive (i.e. the forbidden bits `0137` — group-write, group-execute, and any “other” access — are clear). Because this rec is **Manual**, the script below is an assist that lists files for the operator to review; it does not render a pass/fail verdict.

Run the following script to list any journald log file (referenced by an `f` entry in the `tmpfiles.d` configuration) whose current mode grants any of the forbidden `0137` bits:

Internal Only - General

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

{
 config_files=(/etc/tmpfiles.d/systemd.conf /usr/lib/tmpfiles.d/systemd.conf)
 a_review=()
 for file in "${config_files[@]}"; do
 [ -f "$file" ] || continue
 while read -r type path perm _; do
 if [[ "$type" == "f" && "$perm" =~ ^0?[0-7]{3,4}$ && -f "$path" ]]; then
 actual="$(stat -c '%a' "$path" 2>/dev/null)"
 if [ -n "$actual" ] && [ $(( 0$actual & 0137 )) -ne 0 ]; then
 a_review+=(" - File: \"$path\" | Mode: \"$actual\" | Defined in: \"$file\"
(should be 0640 or more restrictive)")
 fi
 fi
 done < "$file"
 done
 if [ "${#a_review[@]}" -eq 0 ]; then
 printf '%s\n' "No tmpfiles.d-defined journald log files are more permissive
than 0640."
 else
 printf '%s\n' "Files to review:" "${a_review[@]}"
 fi
}
```

Nothing should be returned.

Any file listed is more permissive than `0640` and should be reviewed against site policy. If site policy mandates a less restrictive mode, follow that policy.

Remediation:

If the default configuration is not appropriate for the site specific requirements, copy `/usr/lib/tmpfiles.d/systemd.conf` to `/etc/tmpfiles.d/systemd.conf` and modify as required. Requirements is either `0640` or site policy if that is less restrictive.

References:

1. GRID: NIX-00000266
2. NIST SP 800-53 Rev. 5: AC-3, AU-2, AU-12, MP-2, SI-5

Additional Information:

See `man 5 tmpfiles.d` for detailed information on the permission sets for the relevant log files. Further information with examples can be found at https://www.freedesktop.org/software/systemd/man/tmpfiles.d.html

Internal Only - General

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.	●	●	●

Internal Only - General

6.1.1.1.5 Ensure journald log file rotation is configured (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Journald includes the capability of rotating log files regularly to avoid filling up the system with logs or making the logs unmanageably large. The file `/etc/systemd/journald.conf` is the configuration file used to specify how logs generated by Journald should be rotated.

Rationale:

By keeping the log files smaller and more manageable, a system administrator can easily archive these files to another system and spend less time looking through inordinately large log files.

Impact:

Log rotation limits may cause older entries to be dropped sooner; review the configured thresholds against local retention requirements before applying.

Audit:

Review `/etc/systemd/journald.conf` and files in the `/etc/systemd/journald.conf.d/` directory ending in `.conf`. Verify logs are rotated according to site policy.

Run the following command and ensure logs are rotated according to site policy:

```
# grep -Prsi --
'\b(SystemMaxUse|SystemKeepFree|RuntimeMaxUse|RuntimeKeepFree|MaxFileSec)='
/etc/systemd/journald.conf /etc/systemd/journald.conf.d/
SystemMaxUse=
SystemKeepFree=
RuntimeMaxUse=
RuntimeKeepFree=
MaxFileSec=
```

Remediation:

Edit `/etc/systemd/journald.conf` or a file ending in `.conf` the `/etc/systemd/journald.conf.d/` directory. Set the following parameters in the `[Journal]` section to ensure logs are rotated according to site policy. The settings should be carefully understood as there are specific edge cases and prioritization of parameters.

Example Configuration

Internal Only - General

```
[Journal]
SystemMaxUse=1G
SystemKeepFree=500M
RuntimeMaxUse=200M
RuntimeKeepFree=50M
MaxFileSec=1month
```

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

Run the following command to apply the configuration change:

```
# systemctl reload-or-restart systemd-journald
```

References:

1. GRID: NIX-00000267
2. NIST SP 800-53 Rev. 5: AU-2, AU-7, AU-12
3. https://man7.org/linux/man-pages/man5/journald.conf.5.html

Additional Information:

See `man 5 journald.conf` for detailed information regarding the parameters in use.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 Collect Audit Logs Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.	●	●	●
v7	6.2 Activate audit logging Ensure that local logging has been enabled on all systems and networking devices.	●	●	●
v7	6.3 Enable Detailed Logging Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.		●	●

Internal Only - General

6.1.1.1.6 Ensure journald Storage is configured (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 due to a reboot.

Rationale:

Writing log data to disk will provide the ability to forensically reconstruct events which may have impacted the operations or security of a system even after a system crash or reboot.

Audit:

Run the following script to verify `Storage` is set to `persistent`:

Internal Only - General

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

{
   l_analyze_cmd="$(readlink -e /bin/systemd-analyze || \
   readlink -e /usr/bin/systemd-analyze)"
   l_conf_file="systemd/journald.conf" l_block="Journal"
   l_option="Storage" l_option_value="persistent" a_output=()
   while IFS= read -r l_file; do
      l_file="${l_file//# /}"
      l_opt="$(awk '/\['"$l_block"'\]/{a=1;next}/\[/{a=0}a' "$l_file" \
      2>/dev/null | grep -Poi '^\h*'"$l_option"'\h*=\h*\H+\b' | tail -n 1)"
         l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_option\" is set to: \"$l_option_value\"" \
      "    in: \"$l_file\"")
   done < <("$l_analyze_cmd" cat-config "$l_conf_file" | tac | \
   grep -Pio '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   if [ "${#a_output[@]}" -le "0" ]; then
      l_file="$(readlink -e /etc/"$l_conf_file" || \
      readlink -e /usr/lib/"$l_conf_file")"
      l_opt="$(awk '/\['"$l_block"'\]/{a=1;next}/\[/{a=0}a' "$l_file" \
      2>/dev/null | grep -Poim 1 '^(\h*#)?\h*'"$l_option"'\h*=\h*\H+\b')"
      l_option_value="$(cut -d= -f2 <<< "${l_opt//# /}" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - The default value: \"${l_opt//#/}\"" \
      "    is being used in the configuration")
   fi
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Remediation:

Set the following parameter in the `[Journal]` section in `/etc/systemd/journald.conf` or a file in `/etc/systemd/journald.conf.d/` ending in `.conf`:

```
Storage=persistent
```

Example:

Internal Only - General

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

{
   [ ! -d /etc/systemd/journald.conf.d/ ] && mkdir -p
/etc/systemd/journald.conf.d/
   if grep -Psq -- '^\h*\[Journal\]' /etc/systemd/journald.conf.d/60-
journald.conf; then
      printf '%s\n' "Storage=persistent" >> /etc/systemd/journald.conf.d/60-
journald.conf
   else
      printf '%s\n' "" "[Journal]" "Storage=persistent" >>
/etc/systemd/journald.conf.d/60-journald.conf
   fi
}
```

Note: Drop-in configuration files have higher precedence and override the main configuration file. Files in the `*.conf.d/` configuration subdirectories are sorted by their filename in lexicographic order, regardless of in which of the subdirectories they reside. When multiple files specify the same option, for options which accept just a single value, the entry in the file sorted last takes precedence, and for options which accept a list of values, entries are collected as they occur in the sorted files.

Run to following command to update the parameters in the service:

```
# systemctl reload-or-restart systemd-journald
```

References:

1. GRID: NIX-00000270
2. NIST SP 800-53 Rev. 5: AU-3, AU-12
3. journald.conf(5) - Linux manual page: https://man7.org/linux/man-pages/man5/journald.conf.5.html

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 Collect Audit Logs Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.	●	●	●
v7	6.2 Activate audit logging Ensure that local logging has been enabled on all systems and networking devices.	●	●	●
v7	6.3 Enable Detailed Logging Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.		●	●

Internal Only - General

6.1.1.1.7 Ensure journald Compress is configured (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:

Run the following script to verify `Compress` is set to `yes`:

Internal Only - General

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

{
   l_analyze_cmd="$(readlink -e /bin/systemd-analyze || \
   readlink -e /usr/bin/systemd-analyze)"
   l_conf_file="systemd/journald.conf" l_block="Journal"
   l_option="Compress" l_option_value="yes" a_output=()
   while IFS= read -r l_file; do
      l_file="${l_file//# /}"
      l_opt="$(awk '/\['"$l_block"'\]/{a=1;next}/\[/{a=0}a' "$l_file" \
      2>/dev/null | grep -Poi '^\h*'"$l_option"'\h*=\h*\H+\b' | tail -n 1)"
         l_option_value="$(cut -d= -f2 <<< "$l_opt" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - \"$l_option\" is set to: \"$l_option_value\"" \
      "    in: \"$l_file\"")
   done < <("$l_analyze_cmd" cat-config "$l_conf_file" | tac | \
   grep -Pio '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   if [ "${#a_output[@]}" -le "0" ]; then
      l_file="$(readlink -e /etc/"$l_conf_file" || \
      readlink -e /usr/lib/"$l_conf_file")"
      l_opt="$(awk '/\['"$l_block"'\]/{a=1;next}/\[/{a=0}a' "$l_file" \
      2>/dev/null | grep -Poim 1 '^(\h*#)?\h*'"$l_option"'\h*=\h*\H+\b')"
      l_option_value="$(cut -d= -f2 <<< "${l_opt//# /}" | xargs)"
      [ -n "$l_option_value" ] && \
      a_output+=(" - The default value: \"${l_opt//#/}\"" \
      "    is being used in the configuration")
   fi
   [ "${#a_output[@]}" -gt "0" ] && printf '%s\n' "" "${a_output[@]}" ""
}
```

Remediation:

Set the following parameter in the `[Journal]` section in `/etc/systemd/journald.conf` or a file in `/etc/systemd/journald.conf.d/` ending in `.conf`:

```
Compress=yes
```

Example:

Internal Only - General

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

{
   [ ! -d /etc/systemd/journald.conf.d/ ] && mkdir -p
/etc/systemd/journald.conf.d/
   if grep -Psq -- '^\h*\[Journal\]' /etc/systemd/journald.conf.d/60-
journald.conf; then
      printf '%s\n' "" "Compress=yes" >> /etc/systemd/journald.conf.d/60-
journald.conf
   else
      printf '%s\n' "" "[Journal]" "Compress=yes" >>
/etc/systemd/journald.conf.d/60-journald.conf
   fi
}
```

Note: Drop-in configuration files have higher precedence and override the main configuration file. Files in the `*.conf.d/` configuration subdirectories are sorted by their filename in lexicographic order, regardless of in which of the subdirectories they reside. When multiple files specify the same option, for options which accept just a single value, the entry in the file sorted last takes precedence, and for options which accept a list of values, entries are collected as they occur in the sorted files.

Run to following command to update the parameters in the service:

```
# systemctl reload-or-restart systemd-journald
```

References:

1. GRID: NIX-00000263
2. NIST SP 800-53 Rev. 5: AU-4
3. journald.conf(5) - Linux manual page: https://man7.org/linux/man-pages/man5/journald.conf.5.html

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.	●	●	●
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	6.2 Activate audit logging Ensure that local logging has been enabled on all systems and networking devices.	●	●	●

Internal Only - General

Controls Version	Control	IG 1	IG 2	IG 3
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.		●	●
v7	6.4 Ensure adequate storage for logs Ensure that all systems that store logs have adequate storage space for the logs generated.		●	●

Internal Only - General

6.1.2 Configure rsyslog

The `rsyslog` software package may be used instead of the default `journald` logging mechanism.

Rsyslog has evolved over several decades. For this reason it supports three different configuration formats (“languages”):

- `basic` - previously known as the `sysklogd` format, this is the format best used to express basic things, such as where the statement fits on a single line.
  - It stems back to the original syslog.conf format, in use now for several decades.
  - The most common use case is matching on facility/severity and writing matching messages to a log file.
- `advanced` - previously known as the `RainerScript` format, this format was first available in rsyslog v6 and is the current, best and most precise format for non-trivial use cases where more than one line is needed.
  - Prior to v7, there was a performance impact when using this format that encouraged use of the basic format for best results. Current versions of rsyslog do not suffer from this (historical) performance impact.
  - This new style format is specifically targeted towards more advanced use cases like forwarding to remote hosts that might be partially offline.
- `obsolete legacy` - previously known simply as the `legacy` format, this format is exactly what its name implies: it is obsolete and should not be used when writing new configurations. It was created in the early days (up to rsyslog version 5) where we expected that rsyslog would extend sysklogd just mildly. Consequently, it was primarily aimed at small additions to the original sysklogd format.
  - Practice has shown that it was notoriously hard to use for more advanced use cases, and thus we replaced it with the advanced format.
  - In essence, everything that needs to be written on a single line that starts with a dollar sign is legacy format. Users of this format are encouraged to migrate to the basic or advanced formats.

Note: This section only applies if `rsyslog` is the chosen method for client side logging. Do not apply this section if `journald` is used.

Internal Only - General

6.1.2.1 Ensure rsyslog is installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `rsyslog` software is recommended in environments where `journald` does not meet operation requirements.

Rationale:

The security enhancements of `rsyslog` such as connection-oriented (i.e. TCP) transmission of logs, the option to log to database formats, and the encryption of log data en route to a central logging server) justify installing and configuring the package.

Audit:

- IF - `rsyslog` is being used for logging on the system:

Run the following command to verify `rsyslog` is installed:

```
# dpkg-query -s rsyslog &>/dev/null && echo "rsyslog is installed"
```

Verify the output matches:

```
rsyslog is installed
```

Remediation:

Run the following command to install `rsyslog`:

```
# apt install rsyslog
```

Default Value:

Installed

References:

1. GRID: NIX-00000276
2. NIST SP 800-53 Rev. 5: AU-2, AU-3, AU-12

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 Collect Audit Logs Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.	●	●	●
v7	6.2 Activate audit logging Ensure that local logging has been enabled on all systems and networking devices.	●	●	●
v7	6.3 Enable Detailed Logging Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.		●	●

Internal Only - General

6.1.2.2 Ensure rsyslog service is enabled and active (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Once the `rsyslog` package is installed, ensure that the service is enabled.

Rationale:

If the `rsyslog` service is not enabled to start on boot, the system will not capture logging events.

Audit:

- IF - `rsyslog` is being used for logging on the system:

Run the following command to verify `rsyslog.service` is enabled:

```
# systemctl is-enabled rsyslog.service

enabled
```

Run the following command to verify `rsyslog.service` is active:

```
# systemctl is-active rsyslog.service

active
```

Remediation:

- IF - `rsyslog` is being used for logging on the system:

Run the following commands to unmask, enable, and start `rsyslog.service`:

```
# systemctl unmask rsyslog.service
# systemctl enable rsyslog.service
# systemctl start rsyslog.service
```

References:

1. GRID: NIX-00000280
2. CCI-000366
3. NIST SP 800-53 :: CM-6 b
4. NIST SP 800-53A :: CM-6.1 (iv)
5. STIG ID: RHEL-08-010561 | RULE ID: SV-230298r1017108 | CAT II
6. STIG ID: OL08-00-010561 | RULE ID: SV-248615r991589 | CAT II
7. STIG ID: ALMA-09-053150 | RULE ID: SV-269518r1050401 | CAT II
8. STIG ID: UBTU-22-652010 | RULE ID: SV-260588r991562 | CAT II

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 Collect Audit Logs Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.	●	●	●
v7	6.2 Activate audit logging Ensure that local logging has been enabled on all systems and networking devices.	●	●	●
v7	6.3 Enable Detailed Logging Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.		●	●

Internal Only - General

6.1.2.3 Ensure rsyslog log file creation mode is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`rsyslog` will create logfiles that do not already exist on the system.

The `global()` configuration object `umask`, available in rsyslog 8.26.0+, sets the rsyslogd process' umask. If not specified, the system-provided default is used. The value given must always be a 4-digit octal number, with the initial digit being zero.

The legacy `$umask` parameter sets the `rsyslogd` process' umask. If not specified, the system-provided default is used. The value given must always be a 4-digit octal number, with the initial digit being zero.

The legacy `$FileCreateMode` parameter allows the setting of the mode with which `rsyslogd` creates new files. If not specified, the value `0644` is used. The value given must always be a 4-digit octal number, with the initial digit being zero. Please note that the actual permission depend on `rsyslogd` process `umask`. If in doubt, use `$umask 0000` right at the beginning of the configuration file to remove any restrictions.

The legacy `$FileCreateMode` may be specified multiple times. If so, it specifies the creation mode for all selector lines that follow until the next `$FileCreateMode` parameter. Order of lines is vitally important.

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 to find any configured `$FileCreateMode` directive:

```
# grep -Psi -- '^\h*\$FileCreateMode\h+0[0-7]{3}\b' /etc/rsyslog.conf
/etc/rsyslog.d/*.conf 2>/dev/null
```

Example output:

```
/etc/rsyslog.conf:$FileCreateMode 0640
```

Verify the returned value is `0640` or more restrictive (e.g., `0600`, `0440`, `0400`). If no output is returned, the default mode `0644` applies, which is non-compliant — add a `$FileCreateMode` directive set to `0640` or more restrictive.

Internal Only - General

Remediation:

Edit either `/etc/rsyslog.conf` or a dedicated `.conf` file in `/etc/rsyslog.d/` and set `$FileCreateMode` to `0640` or more restrictive:

```
$FileCreateMode 0640
```

Restart the service:

```
# systemctl restart rsyslog
```

References:

1. GRID: NIX-00000278
2. RSYSLOG.CONF(5)
3. NIST SP 800-53 Rev. 5: AC-3, AC-6, MP-2
4. https://www.rsyslog.com/doc/

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.	●	●	●
v8	8.2 Collect Audit Logs Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.	●	●	●
v7	5.1 Establish Secure Configurations Maintain documented, standard security configuration standards for all authorized operating systems and software.	●	●	●
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.		●	●

Internal Only - General

6.1.2.4 Ensure rsyslog 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 `/etc/rsyslog.conf` and `/etc/rsyslog.d/*.conf` files to ensure appropriate logging is set. In addition, run the following command and verify that the log files are logging information as expected:

```
# ls -l /var/log/
```

Remediation:

Edit the following lines in the `/etc/rsyslog.conf` and `/etc/rsyslog.d/*.conf` files as appropriate for your environment.

> **Note:** The below configuration is shown for example purposes only. Due care should be given to how the organization wishes to store log data.

```
*.emerg                                 :omusrmsg:*
auth,authpriv.*                         /var/log/auth.log
mail.*                                  -/var/log/mail
mail.info                               -/var/log/mail.info
mail.warning                            -/var/log/mail.warn
mail.err                                /var/log/mail.err
cron.*                                  /var/log/cron
*.=warning;*.=err                       -/var/log/warn
*.crit                                  /var/log/warn
*.*;mail.none;news.none                 -/var/log/messages
local0,local1.*                         -/var/log/localmessages
local2,local3.*                         -/var/log/localmessages
local4,local5.*                         -/var/log/localmessages
local6,local7.*                         -/var/log/localmessages
```

Run the following command to reload the `rsyslogd` configuration:

Internal Only - General

```
# systemctl restart rsyslog
```

References:

1. GRID: NIX-00000279
2. See the rsyslog.conf(5) man page for more information.
3. NIST SP 800-53 Rev. 5: AU-2, AU-7, AU-12

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 Collect Audit Logs Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.	●	●	●
v7	6.2 Activate audit logging Ensure that local logging has been enabled on all systems and networking devices.	●	●	●
v7	6.3 Enable Detailed Logging Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.		●	●

Internal Only - General

6.1.2.5 Ensure rsyslog is configured to send logs to a remote log host (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`rsyslog` supports the ability to send log events it gathers to a remote log host or to receive messages from remote hosts, thus enabling centralized log management.

Rationale:

Storing log data on a remote host protects log integrity from local attacks. If an attacker gains root access on the local system, they could tamper with or remove log data that is stored on the local system.

Audit:

Review the `/etc/rsyslog.conf` and `/etc/rsyslog.d/*.conf` files and verify that logs are sent to a central host:

basic format

```
# grep -Psi -- '^\h*\*\.\*\h+@{1,2}\H+' /etc/rsyslog.conf
/etc/rsyslog.d/*.conf 2>/dev/null
```

Output should include `@@<FQDN or IP of remote loghost>`:

Example:

```
*.* @@loghost.example.com
```

advanced format

```
# grep -Psiz -- '^\s*([^#]+\s+)?action\(([^#]+\s+)?\btarget=\"?[^#\"]+\"?\b'
\
  /etc/rsyslog.conf /etc/rsyslog.d/*.conf 2>/dev/null
```

Output should include `target=<FQDN or IP of remote loghost>`:

Example:

Internal Only - General

```
*.* action(type="omfwd" target="loghost.example.com" port="514"
protocol="tcp"
           action.resumeRetryCount="100" queue.type="LinkedList"
queue.size="1000")
```

Remediation:

Edit the `/etc/rsyslog.conf` and `/etc/rsyslog.d/*.conf` files and add the following line (where `loghost.example.com` is the name of your central log host). The `target` directive may either be a fully qualified domain name or an IP address.

Example:

```
*.* action(type="omfwd" target="loghost.example.com" port="514"
protocol="tcp"
           action.resumeRetryCount="100"
           queue.type="LinkedList" queue.size="1000")
```

Run the following command to reload `rsyslog.service`:

```
# systemctl reload-or-restart rsyslog.service
```

References:

1. GRID: NIX-00000275
2. See the rsyslog.conf(5) man page for more information.
3. NIST SP 800-53 Rev. 5: AU-6
4. https://www.rsyslog.com/doc/

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 Collect Audit Logs Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.	●	●	●
v7	6.2 Activate audit logging Ensure that local logging has been enabled on all systems and networking devices.	●	●	●
v7	6.3 Enable Detailed Logging Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.		●	●

Internal Only - General

6.1.2.6 Ensure rsyslog is not configured to receive logs from a remote client (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`rsyslog` supports the ability to receive messages from remote hosts, thus acting as a log server. Clients should not receive data from other hosts.

Rationale:

If a client is configured to also receive data, thus turning it into a server, the client system is acting outside its operational boundary.

Audit:

Review the `/etc/rsyslog.conf` and `/etc/rsyslog.d/*.conf` files and verify that the system is not configured to accept incoming logs.

advanced format — TCP

```
# grep -Psi -- '^\h*module\(load=\"?imtcp\"?\)' /etc/rsyslog.conf
/etc/rsyslog.d/*.conf 2>/dev/null
# grep -Psi -- '^\h*input\(type=\"?imtcp\"?\b' /etc/rsyslog.conf
/etc/rsyslog.d/*.conf 2>/dev/null
```

Nothing should be returned

advanced format — UDP

```
# grep -Psi -- '^\h*module\(load=\"?imudp\"?\)' /etc/rsyslog.conf
/etc/rsyslog.d/*.conf 2>/dev/null
# grep -Psi -- '^\h*input\(type=\"?imudp\"?\b' /etc/rsyslog.conf
/etc/rsyslog.d/*.conf 2>/dev/null
```

Nothing should be returned

obsolete legacy format — TCP

```
# grep -Psi -- '^\h*\$ModLoad\h+imtcp\b' /etc/rsyslog.conf
/etc/rsyslog.d/*.conf 2>/dev/null
# grep -Psi -- '^\h*\$InputTCPServerRun\b' /etc/rsyslog.conf
/etc/rsyslog.d/*.conf 2>/dev/null
```

Nothing should be returned

obsolete legacy format — UDP

Internal Only - General

```
# grep -Psi -- '^\h*\$ModLoad\h+imudp\b' /etc/rsyslog.conf
/etc/rsyslog.d/*.conf 2>/dev/null
# grep -Psi -- '^\h*\$UDPServerRun\b' /etc/rsyslog.conf /etc/rsyslog.d/*.conf
2>/dev/null
```

Nothing should be returned

Remediation:

Should there be any active log server configuration found in the auditing section, modify those files and remove the specific lines highlighted by the audit. Verify none of the following entries are present in any of `/etc/rsyslog.conf` or `/etc/rsyslog.d/*.conf`.

advanced format

```
module(load="imtcp")
input(type="imtcp" port="514")
module(load="imudp")
input(type="imudp" port="514")
```

deprecated legacy format

```
$ModLoad imtcp
$InputTCPServerRun
$ModLoad imudp
$UDPServerRun
```

Restart the service:

```
# systemctl restart rsyslog
```

References:

1. GRID: NIX-00000277
2. NIST SP 800-53 Rev. 5: AU-2, AU-7, AU-12, CM-6
3. https://www.rsyslog.com/doc/index.html

CIS Controls:

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

Internal Only - General

Internal Only - General

6.1.2.7 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/rsyslog` is the configuration file used to rotate log files created by `syslog` or `rsyslog`.

Rationale:

By keeping the log files smaller and more manageable, a system administrator can easily archive these files to another system and spend less time looking through inordinately large log files.

Audit:

Review `/etc/logrotate.conf` and `/etc/logrotate.d/*` and verify logs are rotated according to site policy.

Remediation:

Edit `/etc/logrotate.conf` and `/etc/logrotate.d/*` to ensure logs are rotated according to site policy.

References:

1. GRID: NIX-00000271
2. NIST SP 800-53 Rev. 5: AU-8

Additional Information:

If no `maxage` setting is set for `logrotate` a situation can occur where `logrotate` is interrupted and fails to delete rotated log files. It is recommended to set this to a value greater than the longest any log file should exist on your system to ensure that any such log file is removed but standard rotation settings are not overridden.

Internal Only - General

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.		●	●

Internal Only - General

6.1.2.8 Ensure rsyslog-gnutls is installed (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

TLS protocol support for rsyslog (GnuTLS). This netstream plugin allows rsyslog to send and receive encrypted syslog messages via the syslog-transport-tls IETF standard protocol using GnuTLS.

Note: While this recommendation uses the GnuTLS stream driver (gtls / rsyslog-gnutls package), the OpenSSL stream driver (ossl / rsyslog-openssl package) is also an acceptable alternative. If using OpenSSL, replace DefaultNetstreamDriver gtls with DefaultNetstreamDriver ossl and ensure the rsyslog-openssl package is installed instead of rsyslog-gnutls.

Rationale:

Traditional syslog is a clear-text protocol that means anyone with a sniffer can have a peek at your data. `rsyslog-gnutls` is a secure communications library implementing the SSL, TLS and DTLS protocols and technologies around them, and provides an easy way to encrypt syslog communication.

Impact:

Enabling TLS support using `rsyslog-gnutls` requires certificate management to ensure reliable encrypted logging. If misconfigured, remote log forwarding may fail until valid certificates and trusted CAs are applied.

Audit:

Run the following command and verify `rsyslog-gnutls` is installed:

```
# dpkg-query -s rsyslog-gnutls &>/dev/null && echo "rsyslog-gnutls is
installed"

rsyslog-gnutls is installed
```

Remediation:

Run the following command to install `rsyslog-gnutls`:

```
# apt install rsyslog-gnutls
```

References:

1. GRID: NIX-00000281

Internal Only - General

2. https://www.rsyslog.com/doc/tutorials/tls.html

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 Collect Audit Logs Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.	●	●	●
v7	6.2 Activate audit logging Ensure that local logging has been enabled on all systems and networking devices.	●	●	●
v7	6.3 Enable Detailed Logging Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.		●	●

Internal Only - General

6.1.2.9 Ensure rsyslog forwarding uses gtls (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Network stream drivers are a layer between various parts of `rsyslogd` (e.g. the imtcp module) and the transport layer. They provide sequenced delivery, authentication and confidentiality to the upper layers. Drivers implement different capabilities.

Note: While gtls (GnuTLS) is used in this recommendation, ossl (OpenSSL via rsyslog-openssl) is an acceptable alternative stream driver. Replace DefaultNetstreamDriver gtls with DefaultNetstreamDriver ossl and ensure rsyslog-openssl is installed if using the OpenSSL driver.

Rationale:

`StreamDriver` must be set to `gtls` to enable TLS for encrypting syslog traffic.

Impact:

If the TLS stream driver is not configured, syslog traffic is transmitted in cleartext and is susceptible to interception. Certificate management overhead is required; if the CA certificate expires or the path is misconfigured, encrypted log forwarding will fail until the trust chain is restored.

Audit:

Run the following command and verify that `StreamDriver` is set to `gtls` wherever the netstream driver is configured:

```
# grep -Psi -- '^[^#]*StreamDriver\s*=\s*"?gtls"?' /etc/rsyslog.conf
/etc/rsyslog.d/*.conf 2>/dev/null
```

A match indicates `gtls` is in use (pass). The pattern uses `^[^#]*` to allow `StreamDriver` to appear anywhere on a non-commented line, whether it is set on its own line, inside an `action(...)` block, or in a `global(...)` directive, and whether or not the value is quoted.

Example Output:

```
/etc/rsyslog.d/40-forward.conf: StreamDriver="gtls" StreamDriverMode="1"
StreamDriverAuthMode="anon")
```

Remediation:

Edit `rsyslog.conf` or a `.conf` file in `/etc/rsyslog.d/` to use `gtls`:

Internal Only - General

Example /etc/rsyslog.d/40-forward.conf

```
# certificate files - just CA for a client
global(DefaultNetstreamDriverCAFile="/path/to/contrib/gnutls/ca.pem")

# set up the action for all messages
action(type="omfwd" protocol="tcp" target="s.example.net" port="6514"
       StreamDriver="gtls" StreamDriverMode="1" StreamDriverAuthMode="anon")
```

References:

1. GRID: NIX-00000274
2. https://www.rsyslog.com/doc/tutorials/tls_cert_client.html
3. https://www.rsyslog.com/doc/rainerscript/global.html
4. https://www.rsyslog.com/doc/concepts/netstrm_drvr.html
5. https://www.rsyslog.com/doc/concepts/ns_gtls.html

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.10 Encrypt Sensitive Data in Transit Encrypt sensitive data in transit. Example implementations can include: Transport Layer Security (TLS) and Open Secure Shell (OpenSSH).		●	●
v7	14.4 Encrypt All Sensitive Information in Transit Encrypt all sensitive information in transit.		●	●

Internal Only - General

6.1.2.10 Ensure rsyslog CA certificates are configured (Manual)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

When rsyslog forwards log data to a remote host over TLS, it uses a Certificate Authority (CA) certificate to verify the authenticity of the remote server's TLS certificate. The CA certificate is specified using the `DefaultNetstreamDriverCAFile` global option in rsyslog configuration.

Rationale:

Configuring a trusted CA certificate ensures that rsyslog only connects to log servers whose TLS certificates are signed by a known, trusted authority. Without a properly configured CA certificate, log forwarding may proceed without server identity verification, exposing the system to man-in-the-middle attacks that could intercept, tamper with, or silently discard forwarded log data. This could result in undetected data leakage or loss of audit trail integrity.

Impact:

Proper certificate management is required. If the CA certificate file is missing, expired, or points to an incorrect path, TLS-encrypted log forwarding will fail until the trust chain is restored. Self-signed certificates should not be used as the CA certificate for production environments.

Audit:

Run the following command and verify `rsyslog` is configured to use a CA certificate. Both the global form (`DefaultNetstreamDriverCAFile`) and the per-target form (`StreamDriver.CAFile`) are accepted:

```
# grep -PEsi -- '(DefaultNetstreamDriverCAFile|StreamDriver\.CAFile)\s*='
/etc/rsyslog.conf /etc/rsyslog.d/*.conf 2>/dev/null
```

Example output (global form):

```
/etc/rsyslog.d/40-
forward.conf:DefaultNetstreamDriverCAFile="/etc/ssl/certs/ca-
certificates.crt"
```

Example output (per-target form):

```
/etc/rsyslog.d/40-forward.conf: StreamDriver.CAFile="/etc/ssl/certs/ca-
certificates.crt"
```

Verify each returned path points to a valid, readable CA certificate file:

Internal Only - General

```
# for f in $(grep -Pso '(?<=StreamDriver\.CAFile=")[^"]+' /etc/rsyslog.conf
/etc/rsyslog.d/*.conf 2>/dev/null) \
 $(grep -Pso '(?<=DefaultNetstreamDriverCAFile=")[^"]+' /etc/rsyslog.conf
/etc/rsyslog.d/*.conf 2>/dev/null); do
 [ -n "$f" ] && openssl verify "$f"
done
```

Note: The audit checks only that a CA file is configured (in either form) and that it points to an existing file. It does not validate the certificate chain against the remote server. The CA certificate should come from a known trusted source and must not be self-signed for production log forwarding. Per-target configuration (`action(... StreamDriver.CAFile="...")`) is also a valid configuration starting with rsyslog 8.2108.0 (Debian 12+, Ubuntu 22.04+).

Remediation:

Edit `/etc/rsyslog.conf` or a `.conf` file in `/etc/rsyslog.d/` to configure the correct path to the CA certificate file. Choose one of the three forms below.

Example 1 — global form, system CA bundle:

```
global(DefaultNetstreamDriverCAFile="/etc/ssl/certs/ca-certificates.crt")
```

Example 2 — global form, dedicated CA certificate:

```
global(DefaultNetstreamDriverCAFile="/etc/rsyslog/tls/ca.pem")
```

Example 3 — per-target form (rsyslog 8.2108.0+, Debian 12+ / Ubuntu 22.04+):

```
*.* action(type="omfwd" target="loghost.example.com" port="6514"
protocol="tcp"
 StreamDriver="gtls" StreamDriverMode="1" StreamDriverAuthMode="anon"
 StreamDriver.CAFile="/etc/ssl/certs/ca-certificates.crt"
 action.resumeRetryCount="100"
 queue.type="LinkedList" queue.size="1000")
```

Note: Example 3 is preferred when forwarding to multiple targets with distinct CA chains; each `action(...)` block can carry its own `StreamDriver.CAFile=`. This form requires rsyslog 8.2108.0 or newer.

After editing, restart rsyslog to apply the change:

```
# systemctl restart rsyslog
```

Note: Ensure the chosen CA path is readable under AppArmor. `/etc/ssl/certs/` and `/etc/rsyslog.d/` are allowed by the default Debian rsyslog profile; a custom path (e.g. `/etc/rsyslog/tls/ca.pem`) requires an AppArmor profile change (see Additional Information).

References:

1. GRID: NIX-00000273
2. https://www.rsyslog.com/doc/tutorials/tls_cert_ca.html
3. https://www.rsyslog.com/doc/tutorials/tls.html#client-setup

Internal Only - General

Additional Information:

AppArmor (the default Debian 13 / Ubuntu MAC) confines `rsyslogd` and only permits it to read CA certificate files in profile-allowed paths. The shipped Debian profile allows `/etc/rsyslog.d/` and the system trust store `/etc/ssl/certs/` out of the box. A CA file placed elsewhere (for example a custom `/etc/rsyslog/tls/ca.pem`) will be denied even when ownership/permissions are correct, producing confusing TLS handshake failures rather than an obvious permission error. If a non-default CA path is used, either relocate the CA file under an allowed path or extend the rsyslog AppArmor profile (e.g. a drop-in under `/etc/apparmor.d/local/usr.sbin.rsyslogd` granting `r` on the chosen path) and reload with `apparmor_parser -r`. Check `dmesg` / `journalctl -k | grep -i apparmor` for `DENIED ... rsyslogd ... <capability/file>` entries when diagnosing.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.10 Encrypt Sensitive Data in Transit Encrypt sensitive data in transit. Example implementations can include: Transport Layer Security (TLS) and Open Secure Shell (OpenSSH).		●	●
v8	8.9 Centralize Audit Logs Centralize, to the extent possible, audit log collection and retention across enterprise assets.		●	●
v7	6.5 Central Log Management Ensure that appropriate logs are being aggregated to a central log management system for analysis and review.		●	●
v7	14.4 Encrypt All Sensitive Information in Transit Encrypt all sensitive information in transit.		●	●

Internal Only - General

6.1.3 Configure Logfiles

Internal Only - General

6.1.3.1 Ensure access to all logfiles has been configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Log files contain information from many services on the the local system, or in the event of a centralized log server, others systems logs as well.

In general log files are found in `/var/log/`, although application can be configured to store logs elsewhere. Should your application store logs in another, ensure to run the same test on that location.

The immediate remediation is not durable on its own — files re-created by their service or by log rotation revert to the service default mode; see the "Permanent fix" section of the Remediation for the logrotate `create` and application-umask companion changes.

Rationale:

It is important that log files have the correct permissions to ensure that sensitive data is protected and that only the appropriate users / groups have access to them.

Audit:

Run the following script to verify that files in `/var/log/` have appropriate permissions and ownership:

Internal Only - General

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

{
 a_output=(); a_output2=()
 f_file_test_chk()
 {
 a_out2=()
 maxperm="$( printf '%o' $(( 0777 & ~$perm_mask)) )"
 [ $(( $l_mode & $perm_mask )) -gt 0 ] && \
 a_out2+=(" o Mode: \"$l_mode\" should be \"$maxperm\" or more restrictive")
 [[ ! "$l_user" =~ $l_auser ]] && \
 a_out2+=(" o Owned by: \"$l_user\" and should be owned by \"${l_auser//|/ or
}\"")
 [[ ! "$l_group" =~ $l_agroup ]] && \
 a_out2+=(" o Group owned by: \"$l_group\" and should be group owned by
\"${l_agroup//|/ or }\"")
 [ "${#a_out2[@]}" -gt 0 ] && a_output2+=(" - File: \"$l_fname\" is:"
"${a_out2[@]}")
 }
 while IFS= read -r -d $'\0' l_file; do
 while IFS=: read -r l_fname l_mode l_user l_group; do
 if grep -Pq -- '\/(apt)\h*$' <<< "$(dirname "$l_fname")"; then
 perm_mask='0133' l_auser="root" l_agroup="(root|adm)"; f_file_test_chk
 else
 case "$(basename "$l_fname")" in
 lastlog | lastlog.* | wtmp | wtmp.* | wtmp-* | btmp | btmp.* | btmp-* |
README)
 perm_mask='0113' l_auser="root" l_agroup="(root|utmp)"
 f_file_test_chk ;;
 cloud-init.log* | localmessages* | waagent.log*)
 perm_mask='0133' l_auser="(root|syslog)" l_agroup="(root|adm)"
 f_file_test_chk ;;
 secure{,*.*,.*,-*} | auth.log | syslog | messages)
 perm_mask='0137' l_auser="(root|syslog)" l_agroup="(root|adm)"
 f_file_test_chk ;;
 SSSD | sssd)
 perm_mask='0117' l_auser="(root|SSSD)" l_agroup="(root|SSSD)"
 f_file_test_chk ;;
 gdm | gdm3)
 perm_mask='0117' l_auser="root" l_agroup="(root|gdm|gdm3)"
 f_file_test_chk ;;
 *.journal | *.journal~)
 perm_mask='0137' l_auser="root" l_agroup="(root|systemd-journal)"
 f_file_test_chk ;;
 *)
 perm_mask='0137' l_auser="(root|syslog)" l_agroup="(root|adm)"
 if [ "$l_user" = "root" ] || ! grep -Pq -- "^\h*$(awk -F: '$1=="'"$l_user"'"
{print $7}' /etc/passwd)\b" /etc/shells; then
 ! grep -Pq -- "$l_auser" <<< "$l_user" && l_auser="(root|syslog|$l_user)"
 ! grep -Pq -- "$l_agroup" <<< "$l_group" && l_agroup="(root|adm|$l_group)"
 fi
 f_file_test_chk ;;
 esac
 fi
 done < <(stat -Lc '%n:%#a:%U:%G' "$l_file")
 done < <(find -L /var/log -type f \( -perm /0137 -o ! -user root -o ! -group
root \) -print0)
```


Internal Only - General

```
 if [ "${#a_output2[@]}" -gt 0 ]; then
 printf '%s\n' "${a_output2[@]}"
 fi
}
```

Any file listed in the output is non-compliant and must be corrected. If nothing is returned, all files in `/var/log/` have appropriate permissions and ownership and the system is compliant.

Remediation:

Run the following script to update permissions and ownership on files in `/var/log`.

Although the script is not destructive, ensure that the output is captured in the event that the remediation causes issues.

Internal Only - General

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

{
 a_output2=()
 f_file_test_fix()
 {
 a_out2=()
 maxperm="$( printf '%o' $(( 0777 & ~$perm_mask)) )"
 if [ $(( $l_mode & $perm_mask )) -gt 0 ]; then
 a_out2+=(" o Mode: \"$l_mode\" should be \"$maxperm\" or more restrictive" "
x Removing excess permissions")
 chmod "$l_rperms" "$l_fname"
 fi
 if [[ ! "$l_user" =~ $l_auser ]]; then
 a_out2+=(" o Owned by: \"$l_user\" and should be owned by \"${l_auser//|/ or
}\"" " x Changing ownership to: \"$l_fix_account\"")
 chown "$l_fix_account" "$l_fname"
 fi
 if [[ ! "$l_group" =~ $l_agroup ]]; then
 a_out2+=(" o Group owned by: \"$l_group\" and should be group owned by
\"${l_agroup//|/ or }\"" " x Changing group ownership to:
\"$l_fix_account\"")
 chgrp "$l_fix_account" "$l_fname"
 fi
 [ "${#a_out2[@]}" -gt 0 ] && a_output2+=(" - File: \"$l_fname\" is:"
"${a_out2[@]}")
 }
 l_fix_account='root'
 while IFS= read -r -d $'\0' l_file; do
 while IFS=: read -r l_fname l_mode l_user l_group; do
 if grep -Pq -- '\/(apt)\h*$' <<< "$(dirname "$l_fname")"; then
 perm_mask='0133' l_rperms="u-x,go-wx" l_auser="root" l_agroup="(root|adm)";
f_file_test_fix
 else
 case "$(basename "$l_fname")" in
 lastlog | lastlog.* | wtmp | wtmp.* | wtmp-* | btmp | btmp.* | btmp-* |
README)
 perm_mask='0113' l_rperms="ug-x,o-wx" l_auser="root" l_agroup="(root|utmp)"
 f_file_test_fix ;;
 cloud-init.log* | localmessages* | waagent.log*)
 perm_mask='0133' l_rperms="u-x,go-wx" l_auser="(root|syslog)"
l_agroup="(root|adm)"
 f_file_test_fix ;;
 secure | auth.log | syslog | messages)
 perm_mask='0137' l_rperms="u-x,g-wx,o-rwx" l_auser="(root|syslog)"
l_agroup="(root|adm)"
 f_file_test_fix ;;
 SSSD | sssd)
 perm_mask='0117' l_rperms="ug-x,o-rwx" l_auser="(root|SSSD)"
l_agroup="(root|SSSD)"
 f_file_test_fix ;;
 gdm | gdm3)
 perm_mask='0117' l_rperms="ug-x,o-rwx" l_auser="root"
l_agroup="(root|gdm|gdm3)"
 f_file_test_fix ;;
 *.journal | *.journal~)
```


Internal Only - General

```
 perm_mask='0137' l_rperms="u-x,g-wx,o-rwx" l_auser="root"
l_agroup="(root|systemd-journal)"
 f_file_test_fix ;;
 *)
 perm_mask='0137' l_rperms="u-x,g-wx,o-rwx" l_auser="(root|syslog)"
l_agroup="(root|adm)"
 if [ "$l_user" = "root" ] || ! grep -Pq -- "^\h*$(awk -F: '$1=="'"$l_user"'"
{print $7}' /etc/passwd)\b" /etc/shells; then
 ! grep -Pq -- "$l_auser" <<< "$l_user" && l_auser="(root|syslog|$l_user)"
 ! grep -Pq -- "$l_agroup" <<< "$l_group" && l_agroup="(root|adm|$l_group)"
 fi
 f_file_test_fix ;;
 esac
 fi
 done < <(stat -Lc '%n:%#a:%U:%G' "$l_file")
 done < <(find -L /var/log -type f \( -perm /0137 -o ! -user root -o ! -group
root \) -print0)
 if [ "${#a_output2[@]}" -le 0 ]; then # If all files passed, then we report
no changes
 a_output+=(" - All files in \"/var/log/\" have appropriate permissions and
ownership")
 printf '\n%s' "- All files in \"/var/log/\" have appropriate permissions and
ownership" " o No changes required" ""
 else
 printf '\n%s' "${a_output2[@]}" ""
 fi
}
```

Note: You may also need to change the configuration for your logging software or services for any logs that had incorrect permissions.

If there are services that log to other locations, ensure that those log files have the appropriate access configured.

Permanent fix — preventing re-flagging after log rotation The script above corrects existing files, but log files re-created by their owning service or by log rotation are born with the service's default mode (often `0644`) and will flag again. Make the correct mode durable:

1. **logrotate `create` directive** — for each affected file, edit `/etc/logrotate.d/<service>` and set `create <mode> <user> <group>` so rotated files are created compliant. Example for sysstat — in `/etc/logrotate.d/sysstat`:

```
create 0640 root adm
```

2. **Application umask** — for daemons that write logs outside the logrotate lifecycle, tighten the application's own umask. Example — add near the top of `/usr/libexec/sysstat/sa1`:

```
umask 0027
```

so `sa<DD>` files are created `0640` rather than `0644`.

Internal Only - General

Observed-in-practice on a minimal CIS-hardened Debian 13 install (each with the `/etc/logrotate.d/` file to edit):

- `/var/log/sysstat/sa*` → `/etc/logrotate.d/sysstat` (plus the `sa1` umask above)
- `/var/log/usbguard/usbguard-audit.log*` → `/etc/logrotate.d/usbguard`
- `/var/log/alternatives.log*` → `/etc/logrotate.d/alternatives`

References:

1. GRID: NIX-00000262
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 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.	●	●	●

6.2 System Auditing

The Linux Auditing System operates on a set of rules that collects certain types of system activity to facilitate incident investigation, detect unauthorized access or modification of data. By default events will be logged to `/var/log/audit/audit.log`, which can be configured in `/etc/audit/auditd.conf`.

The following types of audit rules can be specified:

- Control rules: Configuration of the auditing system.
- File system rules: Allow the auditing of access to a particular file or a directory. Also known as file watches.
- System call rules: Allow logging of system calls that any specified program makes.

Audit rules can be set:

Internal Only - General

- On the command line using the `auditctl` utility. These rules are not persistent across reboots.
- In `/etc/audit/audit.rules`. These rules have to be merged and loaded before they are active.

Notes:

- For 64 bit systems that have `arch` as a rule parameter, you will need two rules: one for 64 bit and one for 32 bit systems calls.
- If the auditing system is configured to be locked (`-e 2`), a system reboot will be required in order to load any changes.
- Key names are optional on the rules and will not be used as a compliance auditing. The usage of key names is highly recommended as it facilitates organization and searching, as such, all remediation steps will have key names supplied.
- It is best practice to store the rules, in number prepended files, in `/etc/audit/rules.d/`. Rules must end in a `.rules` suffix. This then requires the use of `augenrules` to merge all the rules into `/etc/audit/audit.rules` based on their their alphabetical (lexical) sort order. All benchmark recommendations follow this best practice for remediation, specifically using the prefix of `50` which is centre weighed if all rule sets make use of the number prepending naming convention.
- Your system may have been customized to change the default `UID_MIN`. All samples output uses `1000`, but this value will not be used in compliance auditing. To confirm the `UID_MIN` for your system, run the following command: `awk '/^\s*UID_MIN/{print $2}' /etc/login.defs`

Normalization

The Audit system normalizes some entries, so when you look at the sample output keep in mind that:

- With regards to users whose login UID is not set, the values `-1` / `unset` / `4294967295` are equivalent and normalized to `-1`.
- When comparing field types and both sides of the comparison is valid fields types, such as`euid!=uid`, then the auditing system may normalize such that the output is `uid!=euid`.
- Some parts of the rule may be rearranged whilst others are dependant on previous syntax. For example, the following two statements are the same:

```
-a always,exit -F arch=b64 -S execve -C uid!=euid -F auid!=-1 -F
key=user_emulation
```

and

Internal Only - General

```
-a always,exit -F arch=b64 -C euid!=uid -F auid!=unset -S execve -k
user_emulation
```

Capacity planning

The recommendations in this section implement auditing policies that not only produces large quantities of logged data, but may also negatively impact system performance. Capacity planning is critical in order not to adversely impact production environments.

- Disk space. If a significantly large set of events are captured, additional on system or off system storage may need to be allocated. If the logs are not sent to a remote log server, ensure that log rotation is implemented else the disk will fill up and the system will halt. Even when logs are sent to a log server, ensure sufficient disk space to allow caching of logs in the case of temporary network outages.
- Disk IO. It is not just the amount of data collected that should be considered, but the rate at which logs are generated.
- CPU overhead. System call rules might incur considerable CPU overhead. Test the systems open/close syscalls per second with and without the rules to gauge the impact of the rules.

Internal Only - General

6.2.1 Configure auditd Service

The capturing of system events provides system administrators with information to allow them to determine if unauthorized access to their system is occurring.

Internal Only - General

6.2.1.1 Ensure auditd packages are installed (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

auditd is the userspace component to the Linux Auditing System. It's responsible for writing audit records to the disk

Rationale:

The capturing of system events provides system administrators with information to allow them to determine if unauthorized access to their system is occurring.

Audit:

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

```
# dpkg-query -s auditd &>/dev/null && echo auditd is installed

auditd is installed
```

Run the following command to verify `audispd-plugins` is installed:

```
# dpkg-query -s audispd-plugins &>/dev/null && echo audispd-plugins is
installed

audispd-plugins is installed
```

Remediation:

Run the following command to Install `auditd` and `audispd-plugins`

```
# apt install auditd audispd-plugins
```

References:

1. GRID: NIX-00000300
2. STIG ID: UBTU-20-010182 | RULE ID: SV-238298r958506 | CAT II
3. STIG ID: UBTU-22-653015 | RULE ID: SV-260591r1015023 | CAT II
4. NIST SP 800-53 Rev. 5: AU-2, AU-3, AU-12, SI-5

Internal Only - General

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.	●	●	●

Internal Only - General

6.2.1.2 Ensure auditd service is enabled and active (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 | grep '^enabled'

enabled
```

Verify result is "enabled".

Run the following command to verify `auditd` is active:

```
# systemctl is-active auditd | grep '^active'

active
```

Verify result is active

Remediation:

Run the following commands to unmask, enable and start `auditd`:

```
# systemctl unmask auditd
# systemctl enable auditd
# systemctl start auditd
```

References:

1. GRID: NIX-00000301
2. NIST SP 800-53 Rev. 5: AU-2, AU-12, SI-5
3. STIG ID: UBTU-20-010182 | RULE ID: SV-238298r958506 | CAT II
4. STIG ID: UBTU-22-653015 | RULE ID: SV-260591r1015023 | CAT II

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 Collect Audit Logs Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.	●	●	●
v7	6.2 Activate audit logging Ensure that local logging has been enabled on all systems and networking devices.	●	●	●
v7	6.3 Enable Detailed Logging Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.		●	●

Internal Only - General

6.2.1.3 Ensure auditing for processes that start prior to auditd is enabled (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Configure `grub2` so that processes that are capable of being audited can be audited even if they start up prior to `auditd` startup.

Rationale:

Audit events need to be captured on processes that start up prior to `auditd` , so that potential malicious activity cannot go undetected.

Audit:

Run the following command:

```
# find /boot -type f -name 'grub.cfg' -exec grep -Ph -- '^\h*linux' {} + |
grep -v 'audit=1'
```

Nothing should be returned.

Remediation:

Edit `/etc/default/grub` or a file in `/etc/default/grub.d/` ending in `.cfg` and append `audit=1` to the existing `GRUB_CMDLINE_LINUX` value.

Note: Files in `/etc/default/grub.d/` are sourced sequentially by `/etc/default/grub`. A bare assignment like `GRUB_CMDLINE_LINUX="audit=1"` will **replace** any existing value (silently dropping parameters such as `quiet splash`). Always expand the existing variable when appending.

Example:

```
# printf '\nGRUB_CMDLINE_LINUX="${GRUB_CMDLINE_LINUX} audit=1"\n' >>
/etc/default/grub.d/40-custom.cfg
```

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

```
# update-grub
```

A reboot is required for the kernel command-line parameter to take effect.

References:

1. GRID: NIX-00000302
2. NIST SP 800-53 Rev. 5: AU-2, AU-3, AU-12

Internal Only - General

3. STIG ID: UBTU-20-010198 | RULE ID: SV-238299r991555 | CAT II
4. STIG ID: UBTU-22-212015 | RULE ID: SV-260471r991555 | CAT II

Additional Information:

This recommendation is designed around the grub2 bootloader, if another bootloader is in use in your environment enact equivalent settings.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 Collect Audit Logs Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.	●	●	●
v7	6.2 Activate audit logging Ensure that local logging has been enabled on all systems and networking devices.	●	●	●

Internal Only - General

6.2.1.4 Ensure audit_backlog_limit is configured (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

In the kernel-level audit subsystem, a socket buffer queue is used to hold audit events. Whenever a new audit event is received, it is logged and prepared to be added to this queue.

The kernel boot parameter `audit_backlog_limit=N`, with `N` representing the amount of messages, will ensure that a queue cannot grow beyond a certain size. If an audit event is logged which would grow the queue beyond this limit, then a failure occurs and is handled according to the system configuration

Rationale:

If an audit event is logged which would grow the queue beyond the `audit_backlog_limit`, then a failure occurs, auditd records will be lost, and potential malicious activity could go undetected.

Audit:

Run the following command and verify the `audit_backlog_limit=` parameter is set:

```
# find /boot -type f -name 'grub.cfg' -exec grep -Ph -- '^\h*linux' {} + |
grep -Pv 'audit_backlog_limit=\d+\b'
```

Nothing should be returned.

Remediation:

Edit `/etc/default/grub` or a file in `/etc/default/grub.d/` ending in `.cfg` and append `audit_backlog_limit=N` to the existing `GRUB_CMDLINE_LINUX` value. The recommended size for `N` is `8192` or larger.

Note: Files in `/etc/default/grub.d/` are sourced sequentially by `/etc/default/grub`. A bare assignment like `GRUB_CMDLINE_LINUX="audit_backlog_limit=8192"` will **replace** any existing value (silently dropping parameters such as `audit=1`, `quiet`, `splash`). Always expand the existing variable when appending.

Example:

Internal Only - General

```
# printf '\nGRUB_CMDLINE_LINUX="${GRUB_CMDLINE_LINUX}
audit_backlog_limit=8192"\n' >> /etc/default/grub.d/40-custom.cfg
```

If you are applying both this recommendation and 6.2.1.3 (`audit=1`), combine both parameters in a single assignment to avoid one overriding the other:

```
# printf '\nGRUB_CMDLINE_LINUX="${GRUB_CMDLINE_LINUX} audit=1
audit_backlog_limit=8192"\n' >> /etc/default/grub.d/40-custom.cfg
```

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

```
# update-grub
```

A reboot is required for the kernel command-line parameter to take effect.

Default Value:

if `audit_backlog_limit` is not set, the system defaults to `audit_backlog_limit=64`

References:

1. GRID: NIX-00000299
2. NIST SP 800-53 Rev. 5: AU-2, AU-3, AU-12

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 Collect Audit Logs Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.	●	●	●
v7	6.2 Activate audit logging Ensure that local logging has been enabled on all systems and networking devices.	●	●	●
v7	6.3 Enable Detailed Logging Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.		●	●

Internal Only - General

6.2.2 Configure Data Retention

When auditing, it is important to carefully configure the storage requirements for audit logs. By default, auditd will max out the log files at 5MB and retain only 4 copies of them. Older versions will be deleted. It is possible on a system that the 20 MBs of audit logs may fill up the system causing loss of audit data. While the recommendations here provide guidance, check your site policy for audit storage requirements.

Internal Only - General

6.2.2.1 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 -Po -- '^\h*max_log_file\h*=\h*\d+\b' /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>
```

Default Value:

`max_log_file = 8`

References:

1. GRID: NIX-00000294
2. NIST SP 800-53 Rev. 5: AU-8

Additional Information:

The `max_log_file` parameter is measured in megabytes.

Other methods of log rotation may be appropriate based on site policy. One example is time-based rotation strategies which don't have native support in auditd configurations. Manual audit of custom configurations should be evaluated for effectiveness and completeness.

Internal Only - General

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.		●	●

Internal Only - General

6.2.2.2 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
```

References:

1. GRID: NIX-00000295
2. NIST SP 800-53 Rev. 5: AU-8

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.		●	●

Internal Only - General

Internal Only - General

6.2.2.3 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 or put the system in single user mode, if no free space is available or an error is detected on the partition that holds the audit log files.

The `disk_full_action` parameter tells the system what action to take when no free space is available on the partition that holds the audit log files. Valid values are `ignore`, `syslog`, `rotate`, `exec`, `suspend`, `single`, and `halt`.

- `ignore`, the audit daemon will issue a syslog message but no other action is taken
- `syslog`, the audit daemon will issue a warning to syslog
- `rotate`, the audit daemon will rotate logs, losing the oldest to free up space
- `exec`, /path-to-script will execute the script. You cannot pass parameters to the script. The script is also responsible for telling the auditd daemon to resume logging once its completed its action
- `suspend`, the audit daemon will stop writing records to the disk
- `single`, the audit daemon will put the computer system in single user mode
- `halt`, the audit daemon will shut down the system

The `disk_error_action` parameter tells the system what action to take when an error is detected on the partition that holds the audit log files. Valid values are `ignore`, `syslog`, `exec`, `suspend`, `single`, and `halt`.

- `ignore`, the audit daemon will not take any action
- `syslog`, the audit daemon will issue no more than 5 consecutive warnings to syslog
- `exec`, /path-to-script will execute the script. You cannot pass parameters to the script
- `suspend`, the audit daemon will stop writing records to the disk
- `single`, the audit daemon will put the computer system in single user mode
- `halt`, the audit daemon will shut down the system

Rationale:

In high security contexts, the risk of detecting unauthorized access or nonrepudiation exceeds the benefit of the system's availability.

Internal Only - General

Impact:

`disk_full_action` parameter:

- Set to `halt` - the `auditd` daemon will shutdown the system when the disk partition containing the audit logs becomes full.
- Set to `single` - the `auditd` daemon will put the computer system in single user mode when the disk partition containing the audit logs becomes full.

`disk_error_action` parameter:

- Set to `halt` - the `auditd` daemon will shutdown the system when an error is detected on the partition that holds the audit log files.
- Set to `single` - the `auditd` daemon will put the computer system in single user mode when an error is detected on the partition that holds the audit log files.
- Set to `syslog` - the `auditd` daemon will issue no more than 5 consecutive warnings to syslog when an error is detected on the partition that holds the audit log files.

Audit:

Run the following command and verify the `disk_full_action` is set to either `halt` or `single`:

```
# grep -Pi -- '^\h*disk_full_action\h*=\h*(halt|single)\b'
/etc/audit/auditd.conf

disk_full_action = <halt|single>
```

Run the following command and verify the `disk_error_action` is set to `syslog`, `single`, or `halt`:

```
# grep -Pi -- '^\h*disk_error_action\h*=\h*(syslog|single|halt)\b'
/etc/audit/auditd.conf

disk_error_action = <syslog|single|halt>
```

Remediation:

Set one of the following parameters in `/etc/audit/auditd.conf` depending on your local security policies.

```
disk_full_action = <halt|single>
disk_error_action = <syslog|single|halt>
```

Example:

Internal Only - General

```
disk_full_action = halt
disk_error_action = halt
```

References:

1. GRID: NIX-00000338
2. NIST SP 800-53 Rev. 5: AU-2, AU-8, AU-12, SI-5
3. AUDITD.CONF(5)

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.		●	●

Internal Only - General

6.2.2.4 Ensure system warns when audit logs are low on space (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

The `auditd` daemon can be configured to halt the system, put the system in single user mode or send a warning message, if the partition that holds the audit log files is low on space.

The `space_left_action` parameter tells the system what action to take when the system has detected that it is starting to get low on disk space. Valid values are `ignore`, `syslog`, `rotate`, `email`, `exec`, `suspend`, `single`, and `halt`.

- `ignore`, the audit daemon does nothing
- `syslog`, the audit daemon will issue a warning to syslog
- `rotate`, the audit daemon will rotate logs, losing the oldest to free up space
- `email`, the audit daemon will send a warning to the email account specified in `action_mail_acct` as well as sending the message to syslog
- `exec`, /path-to-script will execute the script. You cannot pass parameters to the script. The script is also responsible for telling the auditd daemon to resume logging once its completed its action
- `suspend`, the audit daemon will stop writing records to the disk
- `single`, the audit daemon will put the computer system in single user mode
- `halt`, the audit daemon will shut down the system

The `admin_space_left_action` parameter tells the system what action to take when the system has detected that it is low on disk space. Valid values are `ignore`, `syslog`, `rotate`, `email`, `exec`, `suspend`, `single`, and `halt`.

- `ignore`, the audit daemon does nothing
- `syslog`, the audit daemon will issue a warning to syslog
- `rotate`, the audit daemon will rotate logs, losing the oldest to free up space
- `email`, the audit daemon will send a warning to the email account specified in `action_mail_acct` as well as sending the message to syslog
- `exec`, /path-to-script will execute the script. You cannot pass parameters to the script. The script is also responsible for telling the auditd daemon to resume logging once its completed its action
- `suspend`, the audit daemon will stop writing records to the disk
- `single`, the audit daemon will put the computer system in single user mode
- `halt`, the audit daemon will shut down the system

Internal Only - General

Rationale:

In high security contexts, the risk of detecting unauthorized access or nonrepudiation exceeds the benefit of the system's availability.

Impact:

If the `admin_space_left_action` is set to `single` the audit daemon will put the computer system in single user mode.

Audit:

Run the following command and verify the `space_left_action` is set to `email`, `exec`, `single`, or `halt`:

```
grep -Pi -- '^\h*space_left_action\h*=\h*(email|exec|single|halt)\b'
/etc/audit/auditd.conf
```

Verify the output is `email`, `exec`, `single`, or `halt`

Example output

```
space_left_action = email
```

Run the following command and verify the `admin_space_left_action` is set to `single` **- OR -** `halt`:

```
grep -Pi -- '^\h*admin_space_left_action\h*=\h*(single|halt)\b'
/etc/audit/auditd.conf
```

Verify the output is `single` or `halt`

Example output:

```
admin_space_left_action = single
```

Note: A Mail Transfer Agent (MTA) must be installed and configured properly to set `space_left_action = email`

Remediation:

Set the `space_left_action` parameter in `/etc/audit/auditd.conf` to `email`, `exec`, `single`, or `halt`:

Example:

```
space_left_action = email
```

Set the `admin_space_left_action` parameter in `/etc/audit/auditd.conf` to `single` or `halt`:

Example:

```
admin_space_left_action = single
```

Note: A Mail Transfer Agent (MTA) must be installed and configured properly to set `space_left_action = email`

Internal Only - General

References:

1. GRID: NIX-00000339
2. NIST SP 800-53 Rev. 5: AU-2, AU-8, AU-12, SI-5
3. AUDITD.CONF(5)

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.		●	●

Internal Only - General

6.2.3 Configure auditd Rules

The Audit system operates on a set of rules that define what is to be captured in the log files. The following types of Audit rules can be specified:

- **Control rules:** Allow the Audit system's behavior and some of its configuration to be modified.
- **File system rules:** Allow the auditing of access to a particular file or a directory (Also known as file watches).

  The syntax of these watch rules generally follow this format: `-a always,exit -F arch=b64 -F path=path-to-file -F perm=permissions -F key=keyname`

- **System call rules:** Allow logging of system calls that any specified program makes.

  Syscall rules take the general form of: `-a action,list -S syscall -F field=value -k keyname`

Audit rules can be set:

- On the command line using the `auditctl` utility, these rules are not persistent across reboots.
- In a file ending in `.rules` in the `/etc/audit/rules.d/` directory.

Note:

- The rule must start with `-a always,exit`
- The `arch=` field e.g. `arch=b64` is optional but strongly recommended for best performance. `b64` is used in the example, but the arch can also be what is on your system (x86_64, aarch64, ppc64le, s390x, etc).
- The field `-S` all will be added by `auditctl` if it is not included in the rule.
- The rule must include the `path=` field shown in the example output.
- The rule must include the `perm=` field shown in the example output. The options e.g. `wa` may include additional options e.g. `rwxa`
- The `key=` field e.g. `key=scope` is optional but strongly recommended. Though the name scope is used in the example, this can be set to any name.
- The `-w` option e.g. `-w /etc/sudoers -p wa -k scope` is deprecated. While the deprecated `-w` is in a "passing" state, it is recommended to update the rule in accordance with the Remediation Procedure as the `-w` will soon be unsupported.
- File-watching rules that specify `-F perm=` do not require `-S all`; the kernel audit subsystem automatically limits evaluation to relevant syscalls via the filesystem auditing code path. Including `-S all` forces evaluation on every syscall, significantly increasing audit overhead and creating a local denial-of-service attack surface.

Internal Only - General

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

The following script creates and runs a function called `f_audit_rules_format_chk` that may be used to determine if audit rules have deprecated `-w` format or missing `arch=` option:

Internal Only - General

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

f_audit_rules_format_chk()
{
   rules_dir="/etc/audit/rules.d"
   deprecated_flag="-w"
   a_output=()
   for file in "$rules_dir"/*rules; do
   [[ -e "$file" ]] || continue
   while IFS= read -r line; do
   # Skip empty or commented lines
   [[ "$line" =~ ^\s*# ]] || [[ -z "$line" ]] && continue
   # Check for deprecated -w
   if [[ "$line" == *"$deprecated_flag"* ]]; then
   a_output+=("Deprecated '-w' found in $file
   Rule:
   $(wrap_text "$line")
   Flag: $deprecated_flag")
   fi
   # Check for missing -F arch
   if [[ "$line" =~ ^-a[[:space:]]+always,exit ]] && [[ "$line" != *"-F
arch="* ]] \
   ; then
   a_output+=("Missing '-F arch=' in $file
   Rule:
   $(wrap_text "$line")
   Flag: Missing arch")
   fi
   done < "$file"
   done
   if [ ${#a_output[@]} -eq 0 ]; then
   echo -e "No issues found in audit rules files"
   else
   echo -e "\nAudit rules issue detected:"
   for output in "${a_output[@]}"; do
   echo -e "\n$output"
   done
   fi
}
# wrap_text function for formatting
wrap_text() {
    echo "$1"
}
f_audit_rules_format_chk
```


Internal Only - General

6.2.3.1 Ensure modification of the /etc/sudoers file 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`, or files in `/etc/sudoers.d`, will be written to when the file(s) or related attributes have changed. The audit records will be tagged with the identifier "scope".

Rationale:

Changes in the `/etc/sudoers` and `/etc/sudoers.d` files can indicate that an unauthorized change has been made to the scope of system administrator activity.

Satisfies: SRG-OS-000062-GPOS-00031, SRG-OS-000004-GPOS-00004, SRG-OS-000037-GPOS-00015, SRG-OS-000042-GPOS-00020, SRG-OS-000062-GPOS-00031, SRG-OS-000304-GPOS-00121, SRG-OS-000392-GPOS-00172, SRG-OS-000462-GPOS-00206, SRG-OS-000470-GPOS-00214, SRG-OS-000471-GPOS-00215, SRG-OS-000239-GPOS-00089, SRG-OS-000240-GPOS-00090, SRG-OS-000241-GPOS-00091, SRG-OS-000303-GPOS-00120, SRG-OS-000304-GPOS-00121, CCI-002884, SRG-OS-000466-GPOS-00210, SRG-OS-000476-GPOS-00221

Audit:

Note: The `-w` option e.g. `-w /etc/sudoers -p wa -k scope` is deprecated and the rule should be updated in accordance with the Remediation Procedure

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps -- '^\h*[^#\n\r]+\h*\/etc\/sudoers'
```

Verify output includes:

```
-a always,exit -F path=/etc/sudoers -F perm=wa
-a always,exit -F dir=/etc/sudoers.d -F perm=wa
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- '^\h*[^#\n\r]+\h*\/etc\/sudoers' /etc/audit/rules.d/*.rules
```

Example output:

Internal Only - General

```
/etc/audit/rules.d/50-scope.rules:-a always,exit -F arch=b64 -F
path=/etc/sudoers -F perm=wa -k scope
/etc/audit/rules.d/50-scope.rules:-a always,exit -F arch=b64 -F
dir=/etc/sudoers.d -F perm=wa -k scope
```

Remediation:

Note:

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor scope changes for system administrators.

Example:

```
# printf '%s\n' "
-a always,exit -F arch=b64 -F path=/etc/sudoers -F perm=wa -k scope
-a always,exit -F arch=b64 -F dir=/etc/sudoers.d -F perm=wa -k scope
" >> /etc/audit/rules.d/50-scope.rules
```

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000330
2. NIST SP 800-53 :: AU-12 a
3. NIST SP 800-53A :: AU-12.1 (ii)
4. STIG ID: RHEL-08-030171 | RULE ID: SV-230409r627750 | CAT II
5. STIG ID: RHEL-08-030172 | RULE ID: SV-230410r627750 | CAT II
6. STIG ID: RHEL-09-654215 | RULE ID: SV-258217r1045436 | CAT II
7. STIG ID: RHEL-09-654220 | RULE ID: SV-258218r1045439 | CAT II
8. STIG ID: ALMA-09-004970 | RULE ID: SV-269129r1050011 | CAT II
9. STIG ID: ALMA-09-006070 | RULE ID: SV-269135r1050017 | CAT II
10. STIG ID: UBTU-22-654220 | RULE ID: SV-260646r991575 | CAT II
11. STIG ID: UBTU-22-654225 | RULE ID: SV-260647r991575 | CAT II

Internal Only - General

12. STIG ID: UBTU-24-900510 | RULE ID: SV-270807r1066910 | CAT II
13. STIG ID: UBTU-24-900520 | RULE ID: SV-270808r1067100 | CAT II

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	4.8 Log and Alert on Changes to Administrative Group Membership Configure systems to issue a log entry and alert when an account is added to or removed from any group assigned administrative privileges.		●	●

Internal Only - General

6.2.3.2 Ensure actions as another user are always logged (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

`sudo` provides users with temporary elevated privileges to perform operations, either as the superuser or another user.

Rationale:

Creating an audit log of users with temporary elevated privileges and the operation(s) they performed is essential to reporting. Administrators will want to correlate the events written to the audit trail with the records written to `sudo`'s logfile to verify if unauthorized commands have been executed.

Satisfies: SRG-OS-000326-GPOS-00126, SRG-OS-000327-GPOS-00127

Audit:

Note:

- The deprecated `-w` format which allows for backwards compatibility at the expense of system performance is in a "passing" state. However, it is recommended to update the rule in accordance with the Remediation Procedure as the `-w` will soon be unsupported.

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep execve
```

Verify the output includes:

```
-a always,exit -F arch=b64 -S execve -C uid!=euid -F auid!=-1 -F
key=user_emulation
-a always,exit -F arch=b32 -S execve -C uid!=euid -F auid!=-1 -F
key=user_emulation
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- 'execve' /etc/audit/rules.d/*.rules
```

Example output:

Internal Only - General

```
/etc/audit/rules.d/50-user_emulation.rules:-a always,exit -F arch=b32 -C
euid!=uid -F auid!=unset -S execve -k user_emulation
/etc/audit/rules.d/50-user_emulation.rules:-a always,exit -F arch=b64 -C
euid!=uid -F auid!=unset -S execve -k user_emulation
```

Remediation:

Note:

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor elevated privileges.

Example:

```
# printf "
-a always,exit -F arch=b32 -C euid!=uid -F auid!=unset -S execve -k
user_emulation
-a always,exit -F arch=b64 -C euid!=uid -F auid!=unset -S execve -k
user_emulation
" >> /etc/audit/rules.d/50-user_emulation.rules
```

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000286
2. NIST SP 800-53 Revision 5 :: AC-6 (8)
3. https://docs.redhat.com/en/documentation/red_hat_enterprise_linux/10/html/considerations_in_adopting_rhel_10/architectures
4. STIG ID: RHEL-08-030000 | RULE ID: SV-230386r854037 | CAT II
5. STIG ID: RHEL-09-654010 | RULE ID: SV-258176r1045313 | CAT II
6. STIG ID: ALMA-09-007280 | RULE ID: SV-269146r1050028 | CAT II
7. STIG ID: UBTU-22-654230 | RULE ID: SV-260648r958730 | CAT II
8. STIG ID: UBTU-24-200580 | RULE ID: SV-270689r1066556 | CAT II

Internal Only - General

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 Collect Detailed Audit Logs Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.		●	●
v7	4.9 Log and Alert on Unsuccessful Administrative Account Login Configure systems to issue a log entry and alert on unsuccessful logins to an administrative account.		●	●

Internal Only - General

6.2.3.3 Ensure events that modify the sudo log file 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:

Note:

- This recommendation requires that the sudo logfile is configured. For additional information see guidance provided in the recommendation "Ensure sudo log file exists"

1. Run the following command to verify that sudo has a custom log file configured:

```
# grep -rPsi
"^\h*Defaults\h+([^#]+,\h*)?logfile\h*=\h*(\"|\')?\H+(\"|\')?(,\h*\H+\h*)*\h*
(#.*)?$" /etc/sudoers*
```

Example output:

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

2. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps -- 'sudo.log'
```

Verify the output includes:

Internal Only - General

```
-a always,exit -F path=/var/log/sudo.log -F perm=wa
```

3. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- 'sudo.log' /etc/audit/rules.d/*.rules
```

Example output:

```
/etc/audit/rules.d/50-sudo.rules: -a always,exit -F arch=b64 -F
path=/var/log/sudo.log -F perm=wa -k sudo_log_file
```

Remediation:

Note:

- This recommendation requires that the `sudo` logfile is configured. For additional information see guidance provided in the recommendation "Ensure sudo log file exists".
- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

1. 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"
```

2. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor events that modify the sudo log file.

Example:

```
# printf "
-a always,exit -F arch=b64 -F path=/var/log/sudo.log -F perm=wa -k
sudo_log_file
" >> /etc/audit/rules.d/50-sudo.rules
```

Note: `path=` should be set to the same sudo log file location as configured above in step 1.

3. Merge and load the rules into active configuration:

Internal Only - General

```
# augenrules --load
```

4. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000320
2. STIG ID: UBTU-22-654235 | RULE ID: SV-260649r986298 | CAT II
3. STIG ID: UBTU-24-500010 | RULE ID: SV-270740r1066709 | CAT II

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	4.9 Log and Alert on Unsuccessful Administrative Account Login Configure systems to issue a log entry and alert on unsuccessful logins to an administrative account.		●	●

Internal Only - General

6.2.3.4 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
- `clock_settime` - allows for the setting of several internal clocks and timers

system calls have been executed. Further, ensure to write an audit record to the configured audit log file upon exit, tagging the records with a unique identifier such as "time-change".

Rationale:

Unexpected changes in system date and/or time could be a sign of malicious activity on the system.

Audit:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

Run the following command to verify a rule exits in the running configuration:

1a. Verify a rule exists for `adjtimex`, `settimeofday`, `clock_settime` in the running configuration:

```
# auditctl -l | grep -Ps -- '(adjtimex|settimeofday|clock_settime)'
```

Verify output of matches:

```
-a always,exit -S adjtimex,settimeofday
-a always,exit -S clock_settime -F a0=0x0
```

1b. verify a rule for `time-change` exists in the running configuration:

```
# auditctl -l | grep -Ps -- '^\h*[^#\n\r]+\h*\/etc\/localtime'
```

Verify the output matches:

Internal Only - General

```
-a always,exit -F path=/etc/localtime -F perm=wa
```

Run the following command to verify a rule exists in a rules file:

2a. Run the following command to verify a rule exists for `adjtimex`, `settimeofday`, `clock_settime` in a rules file:

```
# grep -Ps -- '(adjtimex|settimeofday|clock_settime)'
/etc/audit/rules.d/*.rules
```

Example output:

```
/etc/audit/rules.d/50-time-change.rules:-a always,exit -F arch=b64 -S
adjtimex,settimeofday -k time-change
/etc/audit/rules.d/50-time-change.rules:-a always,exit -F arch=b64 -S
clock_settime -F a0=0x0 -k time-change
```

2b. Run the following command to verify a rule exists for `/etc/localtime` in a rules file:

```
# grep -Ps -- '^\h*[^#\n\r]+\h*\/etc\/localtime' /etc/audit/rules.d/*.rules
```

Example output:

```
/etc/audit/rules.d/50-time-change.rules:-a always,exit -F arch=b64 -F
path=/etc/localtime -F perm=wa -k localtime-change
```

Remediation:

Note:

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor events that modify date and time information.

1a. Create a rule for `adjtimex`, `settimeofday`, `clock_settime`:

Example

```
# printf "
-a always,exit -F arch=b64 -S adjtimex,settimeofday -k time-change
-a always,exit -F arch=b64 -S clock_settime -F a0=0x0 -k time-change
" >> /etc/audit/rules.d/50-time-change.rules
```

1b. Create a rule for `/etc/localtime`:

Example

Internal Only - General

```
# printf "
-a always,exit -F arch=b64 -F path=/etc/localtime -F perm=wa -k localtime-
change
" >> /etc/audit/rules.d/50-local-time-change.rules
```

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000318
2. NIST SP 800-53 Rev. 5: AU-3, CM-6

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	5.5 Implement Automated Configuration Monitoring Systems 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.		●	●

Internal Only - General

6.2.3.5 Ensure events that modify sethostname and setdomainname are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Record changes to network environment files or system calls. The below parameters monitors the following system calls, and write an audit event on system call exit:

- `sethostname` - set the systems host name
- `setdomainname` - set the systems domain name

Rationale:

Monitoring `sethostname` and `setdomainname` will identify potential unauthorized changes to host and domain name of a system. The changing of these names could potentially break security parameters that are set based on those names.

Audit:

Note:

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps -- 'sethostname|setdomainname'
```

Verify the output matches:

```
-a always,exit -S sethostname,setdomainname
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- 'sethostname|setdomainname' /etc/audit/rules.d/*.rules
```

Example output:

Internal Only - General

```
/etc/audit/rules.d/50-system_locale.rules:-a always,exit -F arch=b32 -S
sethostname,setdomainname -k system-locale
/etc/audit/rules.d/50-system_locale.rules:-a always,exit -F arch=b64 -S
sethostname,setdomainname -k system-locale
```

Remediation:

Note:

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor events that modify the system's network environment.

Example:

```
# printf "
-a always,exit -F arch=b32 -S sethostname,setdomainname -k system-locale
-a always,exit -F arch=b64 -S sethostname,setdomainname -k system-locale
" >> /etc/audit/rules.d/50-system_locale.rules
```

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000319
2. NIST SP 800-53 Rev. 5: AU-3, CM-6

Internal Only - General

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	5.5 Implement Automated Configuration Monitoring Systems 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.		●	●

Internal Only - General

6.2.3.6 Ensure events that modify /etc/issue and /etc/issue.net are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Record changes to network environment files or system calls. The files being monitored are:

- `/etc/issue` and `/etc/issue.net` - messages displayed pre-login

Rationale:

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.

Audit:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps -- '^\h*[^#\n\r]+\h*\/etc\/issue'
```

Verify the output includes:

```
-a always,exit -F path=/etc/issue -F perm=wa
-a always,exit -F path=/etc/issue.net -F perm=wa
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- '^\h*[^#\n\r]+\h*\/etc\/issue' /etc/audit/rules.d/*.rules
```

Example output:

```
/etc/audit/rules.d/50-etc_issue_system_locale.rules:-a always,exit -F
arch=b64 -F path=/etc/issue -F perm=wa -k system-locale
/etc/audit/rules.d/50-etc_issue_system_locale.rules:-a always,exit -F
arch=b64 -F path=/etc/issue.net -F perm=wa -k system-locale
```

Remediation:

Note:

Internal Only - General

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor events that modify the system's network environment.

Example:

```
# printf "
-a always,exit -F arch=b64 -F path=/etc/issue -F perm=wa -k system-locale
-a always,exit -F arch=b64 -F path=/etc/issue.net -F perm=wa -k system-locale
" >> /etc/audit/rules.d/50-etc_issue_system_locale.rules
```

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000310
2. NIST SP 800-53 Rev. 5: AU-3, CM-6

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	5.5 Implement Automated Configuration Monitoring Systems 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.		●	●

Internal Only - General

6.2.3.7 Ensure events that modify /etc/hosts and /etc/hostname are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Record changes to network environment files or system calls. The files being monitored are:

- `/etc/hosts` - file containing host names and associated IP addresses
- `/etc/hostname` - file contains the system's host name

Rationale:

The `/etc/hosts` file is monitored for changes that can indicate an unauthorized intruder is trying to change machine associations with IP addresses and trick users and processes into connecting to unintended machines.

Audit:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps -- '^\h*[^#\n\r]+\h*\/etc\/host'
```

Verify the output includes:

```
-a always,exit -F path=/etc/hosts -F perm=wa
-a always,exit -F path=/etc/hostname -F perm=wa
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- '^\h*[^#\n\r]+\h*\/etc\/host' /etc/audit/rules.d/*.rules
```

Example output:

```
/etc/audit/rules.d/50-etc_host_system_locale.rules:-a always,exit -F arch=b64
-F path=/etc/hosts -F perm=wa -k system-locale
/etc/audit/rules.d/50-etc_host_system_locale.rules:-a always,exit -F arch=b64
-F path=/etc/hostname -F perm=wa -k system-locale
```

Remediation:

Note:

Internal Only - General

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor events that modify the system's network environment.

Example:

```
# printf "
-a always,exit -F arch=b64 -F path=/etc/hosts -F perm=wa -k system-locale
-a always,exit -F arch=b64 -F path=/etc/hostname -F perm=wa -k system-locale
" >> /etc/audit/rules.d/50-etc_host_system_locale.rules
```

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000309
2. NIST SP 800-53 Rev. 5: AU-3, CM-6

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	5.5 Implement Automated Configuration Monitoring Systems 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.		●	●

Internal Only - General

6.2.3.8 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 files being monitored are:

- `/etc/network/interfaces` - This file describes the network interfaces available on your

  system # and how to activate them.

- `/etc/network/interfaces.d/` - directory containing network interface scripts and configurations files
- `/etc/netplan/` - central location for YAML networking configurations files

Rationale:

Monitoring the systems network environment is important as it can show if network interfaces or scripts are being modified in a way that can lead to the machine becoming unavailable or compromised. All audit records should have a relevant tag associated with them.

Audit:

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps -- '\/etc/net'
```

Verify the output includes:

```
-a always,exit -F path=/etc/network/interfaces -F perm=wa
-a always,exit -F dir=/etc/network/interfaces.d -F perm=wa
-a always,exit -F dir=/etc/netplan/ -F perm=wa
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- '\/etc/net' /etc/audit/rules.d/*.rules
```

Example output:

Internal Only - General

```
/etc/audit/rules.d/50-etc_sysconfig_system_locale.rules:-a always,exit -F
arch=b64 -F path=/etc/network/interfaces -F perm=wa -k system-locale
/etc/audit/rules.d/50-etc_sysconfig_system_locale.rules:-a always,exit -F
arch=b64 -F dir=/etc/network/interfaces.d -F perm=wa -k system-locale
/etc/audit/rules.d/50-etc_sysconfig_system_locale.rules:-a always,exit -F
arch=b64 -F dir=/etc/netplan/ -F perm=wa -k system-locale
```

Note: If the directory does not exist by default, this is considered a passing state. Artifacts contain existence checks within the sce.

Remediation:

Note:

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor events that modify the system's network environment.

Example:

```
# printf "
-a always,exit -F arch=b64 -F path=/etc/network/interfaces -F perm=wa -k
system-locale
-a always,exit -F arch=b64 -F dir=/etc/network/interfaces.d -F perm=wa -k
system-locale
-a always,exit -F arch=b64 -F dir=/etc/netplan/ -F perm=wa -k system-locale
" >> /etc/audit/rules.d/50-etc_sysconfig_system_locale.rules
```

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000322
2. NIST SP 800-53 Rev. 5: AU-3, CM-6

Internal Only - General

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	5.5 Implement Automated Configuration Monitoring Systems 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.		●	●

Internal Only - General

6.2.3.9 Ensure events that modify /etc/NetworkManager directory are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Record changes to network environment files or system calls. The files being monitored are:

- `/etc/NetworkManager/` - directory contains configuration files and settings used by the `NetworkManager`

Rationale:

Monitoring the `/etc/NetworkManager/` directory will identify potential unauthorized changes to network configuration files.

Audit:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps -- '\/etc/NetworkManager'
```

Verify the output includes:

```
-a always,exit -F dir=/etc/NetworkManager/ -F perm=wa
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- '\/etc\/NetworkManager' /etc/audit/rules.d/*.rules
```

Example output:

```
/etc/audit/rules.d/50-etc_NetworkManager_system_locale.rules:-a always,exit -
F arch=b64 -F dir=/etc/NetworkManager/ -F perm=wa -k system-locale
```

Note: NetworkManager is not installed by default on Debian 13, this is considered a **PASSING** state. The artifact sce contain existence checks.

Remediation:

Note:

Internal Only - General

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor events that modify the system's network environment.

Example:

```
# printf "
-a always,exit -F arch=b64 -F dir=/etc/NetworkManager/ -F perm=wa -k system-
locale
" >> /etc/audit/rules.d/50-etc_NetworkManager_system_locale.rules
```

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000311
2. NIST SP 800-53 Rev. 5: AU-3, CM-6

Additional Information:

Ticket #27388: Remove -S all from file-watching audit rules (-F perm= present). Redundant flag is a local DoS attack surface. Adds rationale note.

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.		●	●

Internal Only - General

Controls Version	Control	IG 1	IG 2	IG 3
v7	5.5 Implement Automated Configuration Monitoring Systems 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.		●	●

Internal Only - General

6.2.3.10 Ensure use of privileged commands are 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.

Impact:

Both the audit and remediation section of this recommendation will traverse all mounted file systems that is not mounted with either `noexec` or `nosuid` mount options. If there are large file systems without these mount options, **such traversal will be significantly detrimental to the performance of the system.**

Before running either the audit or remediation section, inspect the output of the following command to determine exactly which file systems will be traversed:

```
# findmnt -n -l -k -it $(awk '/nodev/ { print $2 }' /proc/filesystems | paste
-sd,) | grep -Pv "noexec|nosuid"
```

To exclude a particular file system due to adverse performance impacts, update the audit and remediation sections by adding a sufficiently unique string to the `grep` statement. The above command can be used to test the modified exclusions.

Audit:

On disk configuration

Run the following script to check on disk rules:

Internal Only - General

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

{
   for PARTITION in $(findmnt -n -l -k -it $(awk '/nodev/ { print $2 }'
/proc/filesystems | paste -sd,) | grep -Pv "noexec|nosuid" | awk '{print
$1}'); do
      for PRIVILEGED in $(find "${PARTITION}" -xdev -perm /6000 -type f); do
         grep -qr "${PRIVILEGED}" /etc/audit/rules.d && printf "OK:
'${PRIVILEGED}' found in auditing rules.\n" || printf "Warning:
'${PRIVILEGED}' not found in on disk configuration.\n"
      done
   done
}
```

Verify that all output is `OK`.

Running configuration

Run the following script to check loaded rules:

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

{
   RUNNING=$(auditctl -l)
   [ -n "${RUNNING}" ] && for PARTITION in $(findmnt -n -l -k -it $(awk
'/nodev/ { print $2 }' /proc/filesystems | paste -sd,) | grep -Pv
"noexec|nosuid" | awk '{print $1}'); do
      for PRIVILEGED in $(find "${PARTITION}" -xdev -perm /6000 -type f); do
         printf -- "${RUNNING}" | grep -q "${PRIVILEGED}" && printf "OK:
'${PRIVILEGED}' found in auditing rules.\n" || printf "Warning:
'${PRIVILEGED}' not found in running configuration.\n"
      done
   done \
   || printf "ERROR: Variable 'RUNNING' is unset.\n"
}
```

Verify that all output is `OK`.

Special mount points

If there are any special mount points that are not visible by default from `findmnt` as per the above audit, said file systems would have to be manually audited.

Remediation:

Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor the use of privileged commands.

Example script:

Internal Only - General

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

{
  UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
  AUDIT_RULE_FILE="/etc/audit/rules.d/50-privileged.rules"
  NEW_DATA=()
  for PARTITION in $(findmnt -n -l -k -it $(awk '/nodev/ { print $2 }'
/proc/filesystems | paste -sd,) | grep -Pv "noexec|nosuid" | awk '{print
$1}'); do
    readarray -t DATA < <(find "${PARTITION}" -xdev -perm /6000 -type f | awk
-v UID_MIN=${UID_MIN} '{print "-a always,exit -F arch=b64 -S all -F path=" $1
" -F perm=x -F auid>="UID_MIN" -F auid!=unset -k privileged" }')
      for ENTRY in "${DATA[@]}"; do
        NEW_DATA+=("${ENTRY}")
      done
  done
  readarray &> /dev/null -t OLD_DATA < "${AUDIT_RULE_FILE}"
  COMBINED_DATA=( "${OLD_DATA[@]}" "${NEW_DATA[@]}" )
  printf '%s\n' "${COMBINED_DATA[@]}" | sort -u > "${AUDIT_RULE_FILE}"
}
```

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required.

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

Special mount points

If there are any special mount points that are not visible by default from just scanning `/`, change the `PARTITION` variable to the appropriate partition and re-run the remediation.

References:

1. GRID: NIX-00000346
2. NIST SP 800-53 Rev. 5: AU-3, AU-3(1)

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

Internal Only - General

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.	●	●	●

Internal Only - General

6.2.3.11 Ensure unsuccessful file access attempts are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Monitor for unsuccessful attempts to access files. The following parameters are associated with system calls that control files:

- creation - `creat`
- opening - `open` , `openat`
- truncation - `truncate` , `ftruncate`

An audit log record will only be written if all of the following criteria is met for the user when trying to access a file:

- a non-privileged user (`auid>=``UID_MIN`)
- is not a Daemon event (`auid=``4294967295/unset/-1`)
- if the system call returned `EACCES` (permission denied) or EPERM (some other permanent error associated with the specific system call)

Rationale:

Failed attempts to open, create or truncate files could be an indication that an individual or process is trying to gain unauthorized access to the system.

Satisfies: SRG-OS-000062-GPOS-00031, SRG-OS-000037-GPOS-00015, SRG-OS-000042-GPOS-00020, SRG-OS-000392-GPOS-00172, SRG-OS-000462-GPOS-00206, SRG-OS-000471-GPOS-00215, SRG-OS-000064-GPOS-00033

Audit:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Run the following command to verify `UID_MIN` is set in `/etc/login.defs`

```
# awk '/^\s*UID_MIN/{print $2}' /etc/login.defs
```

Example output:

```
1000
```

2. Run the following command to verify a rule exists in the running configuration:

Internal Only - General

```
# auditctl -l | grep -Ps -- '(EACCES|EPERM)'
```

Verify the output matches:

```
-a always,exit -S creat,open,openat,truncate,ftruncate -F exit=-EACCES -F
auid>=1000 -F auid!=unset
-a always,exit -S creat,open,openat,truncate,ftruncate -F exit=-EPERM -F
auid>=1000 -F auid!=unset
```

3. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- '(EACCES|EPERM)' /etc/audit/rules.d/*.rules
```

Example output:

```
/etc/audit/rules.d/50-access.rules:-a always,exit -F arch=b32 -S
creat,open,openat,truncate,ftruncate -F exit=-EACCES -F auid>=1000 -F
auid!=unset -k access
/etc/audit/rules.d/50-access.rules:-a always,exit -F arch=b32 -S
creat,open,openat,truncate,ftruncate -F exit=-EPERM -F auid>=1000 -F
auid!=unset -k access
/etc/audit/rules.d/50-access.rules:-a always,exit -F arch=b64 -S
creat,open,openat,truncate,ftruncate -F exit=-EACCES -F auid>=1000 -F
auid!=unset -k access
/etc/audit/rules.d/50-access.rules:-a always,exit -F arch=b64 -S
creat,open,openat,truncate,ftruncate -F exit=-EPERM -F auid>=1000 -F
auid!=unset -k access
```

Remediation:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Edit the `/etc/login.defs` file and create or add the `UID_MIN` setting.

Example

```
UID_MIN 1000
```

2. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor unsuccessful file access attempts.

Example:

Internal Only - General

```
# printf "
/etc/audit/rules.d/50-access.rules:-a always,exit -F arch=b32 -S
creat,open,openat,truncate,ftruncate -F exit=-EACCES -F auid>=1000 -F
auid!=unset -k access
/etc/audit/rules.d/50-access.rules:-a always,exit -F arch=b32 -S
creat,open,openat,truncate,ftruncate -F exit=-EPERM -F auid>=1000 -F
auid!=unset -k access
-a always,exit -F arch=b64 -S creat,open,openat,truncate,ftruncate -F exit=-
EACCES -F auid>=1000 -F auid!=unset -k access
-a always,exit -F arch=b64 -S creat,open,openat,truncate,ftruncate -F exit=-
EPERM -F auid>=1000 -F auid!=unset -k access
" >> /etc/audit/rules.d/50-access.rules
```

Note: The `auid>=` value should be set to the same value as `UID_MIN` value in `/etc/login.defs`.

3. Merge and load the rules into active configuration:

```
# augenrules --load
```

4. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000345
2. NIST SP 800-53 :: AU-12 a
3. NIST SP 800-53A :: AU-12.1 (ii)
4. STIG ID: RHEL-08-030420 | RULE ID: SV-230449r810455 | CAT II
5. STIG ID: RHEL-09-654070 | RULE ID: SV-258188r1045349 | CAT II
6. STIG ID: ALMA-09-048090 | RULE ID: SV-269475r1050358 | CAT II
7. STIG ID: UBTU-22-654165 | RULE ID: SV-260635r958446 | CAT II
8. STIG ID: UBTU-24-900160 | RULE ID: SV-270787r1068378 | CAT II

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 Collect Detailed Audit Logs Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.		●	●
v7	14.9 Enforce Detail Logging for Access or Changes to Sensitive Data Enforce detailed audit logging for access to sensitive data or changes to sensitive data (utilizing tools such as File Integrity Monitoring or Security Information and Event Monitoring).			●

Internal Only - General

6.2.3.12 Ensure events that modify /etc/group information are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Record events affecting the modification of user or group information, including that of passwords and old passwords if in use.

- `/etc/group` - system groups

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.

Satisfies: SRG-OS-000062-GPOS-00031, SRG-OS-000004-GPOS-00004, SRG-OS-000037-GPOS-00015, SRG-OS-000042-GPOS-00020, SRG-OS-000062-GPOS-00031, SRG-OS-000304-GPOS-00121, SRG-OS-000392-GPOS-00172, SRG-OS-000462-GPOS-00206, SRG-OS-000470-GPOS-00214, SRG-OS-000471-GPOS-00215, SRG-OS-000239-GPOS-00089, SRG-OS-000240-GPOS-00090, SRG-OS-000241-GPOS-00091, SRG-OS-000303-GPOS-00120, SRG-OS-000304-GPOS-00121, SRG-OS-000476-GPOS-00221

Audit:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps -- '^\h*[^#\n\r]+\h*\/etc\/group'
```

Verify the output includes:

```
-a always,exit -F path=/etc/group -F perm=wa
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- '^\h*[^#\n\r]+\h*\/etc\/group' /etc/audit/rules.d/*.rules
```

Example output:

Internal Only - General

```
/etc/audit/rules.d/50-identity.rules:-a always,exit -F arch=b64 -F
path=/etc/group -F perm=wa -k identity
```

Remediation:

Note:

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor events that modify user/group information.

Example:

```
# printf "
-a always,exit -F arch=b64 -F path=/etc/group -F perm=wa -k identity
" >> /etc/audit/rules.d/50-identity.rules
```

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000308
2. https://manpages.debian.org/bookworm/manpages/nsswitch.conf.5.en.html
3. https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/6/html/managing_smart_cards/pam_configuration_files
4. NIST SP 800-53 :: AU-12 a
5. NIST SP 800-53A :: AU-12.1 (ii)
6. STIG ID: RHEL-08-030170 | RULE ID: SV-230408r1017214 | CAT II
7. STIG ID: RHEL-09-654225 | RULE ID: SV-258227r1014992 | CAT II
8. STIG ID: ALMA-09-005080 | RULE ID: SV-269129r1050011 | CAT II
9. STIG ID: UBTU-22-654130 | RULE ID: SV-260628r958368 | CAT II
10. STIG ID: UBTU-24-200290 | RULE ID: SV-270685r1066544 | CAT II

Internal Only - General

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	4.8 Log and Alert on Changes to Administrative Group Membership Configure systems to issue a log entry and alert when an account is added to or removed from any group assigned administrative privileges.		●	●

Internal Only - General

6.2.3.13 Ensure events that modify /etc/passwd information are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Record events affecting the modification of user or group information, including that of passwords and old passwords if in use.

- `/etc/passwd` - system users

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.

Satisfies: SRG-OS-000062-GPOS-00031, SRG-OS-000004-GPOS-00004, SRG-OS-000037-GPOS-00015, SRG-OS-000042-GPOS-00020, SRG-OS-000062-GPOS-00031, SRG-OS-000304-GPOS-00121, SRG-OS-000392-GPOS-00172, SRG-OS-000462-GPOS-00206, SRG-OS-000470-GPOS-00214, SRG-OS-000471-GPOS-00215, SRG-OS-000239-GPOS-00089, SRG-OS-000240-GPOS-00090, SRG-OS-000241-GPOS-00091, SRG-OS-000303-GPOS-00120, SRG-OS-000304-GPOS-00121, SRG-OS-000476-GPOS-00221

Audit:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps -- '^\h*[^#\n\r]+\h*\/etc\/passwd'
```

Verify the output includes:

```
-a always,exit -F path=/etc/passwd -F perm=wa
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- '^\h*[^#\n\r]+\h*\/etc\/passwd' /etc/audit/rules.d/*.rules
```

Example output:

Internal Only - General

```
/etc/audit/rules.d/50-identity.rules:-a always,exit -F arch=b64 -F
path=/etc/passwd -F perm=wa -k identity
```

Remediation:

Note:

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor events that modify user/group information.

Example:

```
# printf "
-a always,exit -F arch=b64 -F path=/etc/passwd -F perm=wa -k identity
" >> /etc/audit/rules.d/50-identity.rules
```

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000314
2. https://manpages.debian.org/bookworm/manpages/nsswitch.conf.5.en.html
3. https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/6/html/managing_smart_cards/pam_configuration_files
4. NIST SP 800-53 :: AU-12 a
5. NIST SP 800-53A :: AU-12.1 (ii)
6. STIG ID: RHEL-08-030150 | RULE ID: SV-230406r1017212 | CAT II
7. STIG ID: RHEL-09-654240 | RULE ID: SV-258230r958408 | CAT II
8. STIG ID: ALMA-09-005410 | RULE ID: SV-269133r1050015 | CAT II
9. STIG ID: UBTU-22-654145 | RULE ID:SV-260631r958368 | CAT II
10. STIG ID: UBTU-24-200280| RULE ID: SV-270684r1066541 | CAT II

Internal Only - General

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	4.8 Log and Alert on Changes to Administrative Group Membership Configure systems to issue a log entry and alert when an account is added to or removed from any group assigned administrative privileges.		●	●

Internal Only - General

6.2.3.14 Ensure events that modify /etc/shadow and /etc/gshadow are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Record events affecting the modification of user or group information, including that of passwords and old passwords if in use.

- `/etc/gshadow` - encrypted password for each group
- `/etc/shadow` - system user passwords

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.

Satisfies: SRG-OS-000062-GPOS-00031, SRG-OS-000004-GPOS-00004, SRG-OS-000037-GPOS-00015, SRG-OS-000042-GPOS-00020, SRG-OS-000062-GPOS-00031, SRG-OS-000304-GPOS-00121, SRG-OS-000392-GPOS-00172, SRG-OS-000462-GPOS-00206, SRG-OS-000470-GPOS-00214, SRG-OS-000471-GPOS-00215, SRG-OS-000239-GPOS-00089, SRG-OS-000240-GPOS-00090, SRG-OS-000241-GPOS-00091, SRG-OS-000303-GPOS-00120, SRG-OS-000304-GPOS-00121, SRG-OS-000476-GPOS-00221

Audit:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps -- '^\h*[^#\n\r]+\h*shadow'
```

Verify the output includes:

```
-a always,exit -F path=/etc/gshadow -F perm=wa
-a always,exit -F path=/etc/shadow -F perm=wa
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- 'shadow' /etc/audit/rules.d/*.rules
```

Example output:

Internal Only - General

```
/etc/audit/rules.d/50-identity.rules:-a always,exit -F arch=b64 -F
path=/etc/gshadow -F perm=wa -k identity
/etc/audit/rules.d/50-identity.rules:-a always,exit -F arch=b64 -F
path=/etc/shadow -F perm=wa -k identity
```

Remediation:

Note:

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor events that modify user/group information.

Example:

```
# printf "
-a always,exit -F arch=b64 -F path=/etc/gshadow -F perm=wa -k identity
-a always,exit -F arch=b64 -F path=/etc/shadow -F perm=wa -k identity
" >> /etc/audit/rules.d/50-identity.rules
```

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000316
2. https://manpages.debian.org/bookworm/manpages/nsswitch.conf.5.en.html
3. https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/6/html/managing_smart_cards/pam_configuration_files
4. NIST SP 800-53 :: AU-12 a
5. NIST SP 800-53A :: AU-12.1 (ii)
6. STIG ID: RHEL-08-030130 | RULE ID: SV-230404r1017210 | CAT II
7. STIG ID: RHEL-09-654245 | RULE ID: SV-258231r1069375 | CAT II
8. STIG ID: ALMA-09-005960 | RULE ID: SV-269134r1050016 | CAT II
9. STIG ID: UBTU-22-654150 | RULE ID: SV-260632r958368 | CAT II

Internal Only - General

10. STIG ID: UBTU-24-200300 | RULE ID: SV-270686r1066547 | CAT II
11. STIG ID: RHEL-08-030160 | RULE ID: SV-230407r1017213 | CAT II
12. STIG ID: RHEL-09-654230 | RULE ID: SV-258228r991572 | CAT II
13. STIG ID: ALMA-09-005190 | RULE ID: SV-269131r1050013 | CAT II
14. STIG ID: UBTU-22-654135 | RULE ID: SV-260629r958368 | CAT II
15. STIG ID: UBTU-24-200310 | RULE ID: SV-270687r1066550 | CAT II

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	4.8 Log and Alert on Changes to Administrative Group Membership Configure systems to issue a log entry and alert when an account is added to or removed from any group assigned administrative privileges.		●	●

Internal Only - General

6.2.3.15 Ensure events that modify /etc/security/opasswd are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Record events affecting the modification of user or group information, including that of passwords and old passwords if in use.

- `/etc/security/opasswd` - storage of old passwords if the relevant PAM module is in use

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.

Satisfies: SRG-OS-000062-GPOS-00031, SRG-OS-000004-GPOS-00004, SRG-OS-000037-GPOS-00015, SRG-OS-000042-GPOS-00020, SRG-OS-000062-GPOS-00031, SRG-OS-000304-GPOS-00121, SRG-OS-000392-GPOS-00172, SRG-OS-000462-GPOS-00206, SRG-OS-000470-GPOS-00214, SRG-OS-000471-GPOS-00215, SRG-OS-000239-GPOS-00089, SRG-OS-000240-GPOS-00090, SRG-OS-000241-GPOS-00091, SRG-OS-000303-GPOS-00120, SRG-OS-000304-GPOS-00121, SRG-OS-000476-GPOS-00221

Audit:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps -- '^\h*[^#\n\r]+\h*\/etc\/security\/opasswd'
```

Verify the output includes:

```
-a always,exit -F path=/etc/security/opasswd -F perm=wa
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- '^\h*[^#\n\r]+\h*\/etc\/security\/opasswd'
/etc/audit/rules.d/*.rules
```

Example output:

Internal Only - General

```
/etc/audit/rules.d/50-identity.rules:-a always,exit -F arch=b64 -F
path=/etc/security/opasswd -F perm=wa -k identity
```

Remediation:

Note:

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor events that modify user/group information.

Example:

```
# printf "
-a always,exit -F arch=b64 -F path=/etc/security/opasswd -F perm=wa -k
identity
" >> /etc/audit/rules.d/50-identity.rules
```

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000315
2. https://manpages.debian.org/bookworm/manpages/nsswitch.conf.5.en.html
3. https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/6/html/managing_smart_cards/pam_configuration_files
4. NIST SP 800-53 :: AU-12 a
5. NIST SP 800-53A :: AU-12.1 (ii)
6. STIG ID: RHEL-08-030140| RULE ID: SV-230405r1017211 | CAT II
7. STIIG ID: RHEL-09-654235 | RULE ID: SV-258229r958434 | CAT II
8. STIG ID: ALMA-09-005300 | RULE ID: SV-269132r1050014 | CAT II
9. STIG ID: UBTU-22-654140 | RULE ID: SV-260630r958368 | CAT II
10. STIG ID: UBTU-24-200320 | RULE ID: SV-270688r1066553 | CAT II

Internal Only - General

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	4.8 Log and Alert on Changes to Administrative Group Membership Configure systems to issue a log entry and alert when an account is added to or removed from any group assigned administrative privileges.		●	●

Internal Only - General

6.2.3.16 Ensure events that modify /etc/nsswitch.conf file are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Record events affecting the modification of user or group information, including that of passwords and old passwords if in use.

- `/etc/nsswitch.conf` - file configures how the system uses various databases and name resolution mechanisms

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.

Satisfies: SRG-OS-000062-GPOS-00031, SRG-OS-000004-GPOS-00004, SRG-OS-000037-GPOS-00015, SRG-OS-000042-GPOS-00020, SRG-OS-000062-GPOS-00031, SRG-OS-000304-GPOS-00121, SRG-OS-000392-GPOS-00172, SRG-OS-000462-GPOS-00206, SRG-OS-000470-GPOS-00214, SRG-OS-000471-GPOS-00215, SRG-OS-000239-GPOS-00089, SRG-OS-000240-GPOS-00090, SRG-OS-000241-GPOS-00091, SRG-OS-000303-GPOS-00120, SRG-OS-000304-GPOS-00121, SRG-OS-000476-GPOS-00221

Audit:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps -- '^\h*[^#\n\r]+\h*\/etc\/nsswitch.conf'
```

Verify the output includes:

```
-a always,exit -F path=/etc/nsswitch.conf -F perm=wa
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- 'nsswitch.conf' /etc/audit/rules.d/*.rules
```

Example output:

Internal Only - General

```
/etc/audit/rules.d/50-identity.rules:-a always,exit -F arch=b64 -F
path=/etc/nsswitch.conf -F perm=wa -k identity
```

Remediation:

Note:

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor events that modify user/group information.

Example:

```
# printf "
-a always,exit -F arch=b64 -F path=/etc/nsswitch.conf -F perm=wa -k identity
" >> /etc/audit/rules.d/50-identity.rules
```

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000312
2. https://manpages.debian.org/bookworm/manpages/nsswitch.conf.5.en.html
3. https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/6/html/managing_smart_cards/pam_configuration_files
4. NIST SP 800-53 :: AU-12 a
5. NIST SP 800-53A :: AU-12.1 (ii)

Internal Only - General

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	4.8 Log and Alert on Changes to Administrative Group Membership Configure systems to issue a log entry and alert when an account is added to or removed from any group assigned administrative privileges.		●	●

Internal Only - General

6.2.3.17 Ensure events that modify /etc/pam.conf and /etc/pam.d/ information are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Record events affecting the modification of user or group information, including that of passwords and old passwords if in use.

- `/etc/pam.conf` - file determines the authentication services to be used, and the order in which the services are used.
- `/etc/pam.d` - directory contains the PAM configuration files for each PAM-aware application.

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.

Satisfies: SRG-OS-000062-GPOS-00031, SRG-OS-000004-GPOS-00004, SRG-OS-000037-GPOS-00015, SRG-OS-000042-GPOS-00020, SRG-OS-000062-GPOS-00031, SRG-OS-000304-GPOS-00121, SRG-OS-000392-GPOS-00172, SRG-OS-000462-GPOS-00206, SRG-OS-000470-GPOS-00214, SRG-OS-000471-GPOS-00215, SRG-OS-000239-GPOS-00089, SRG-OS-000240-GPOS-00090, SRG-OS-000241-GPOS-00091, SRG-OS-000303-GPOS-00120, SRG-OS-000304-GPOS-00121, SRG-OS-000476-GPOS-00221

Audit:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps -- '^\h*[^#\n\r]+\h*\/etc\/pam'
```

Verify the output includes:

```
-a always,exit -F path=/etc/pam.conf -F perm=wa
-a always,exit -F dir=/etc/pam.d -F perm=wa
```

2. Run the following command to verify a rule exists in a rules file:

Internal Only - General

```
# grep -Ps -- '^\h*[^#\n\r]+\h*\/etc\/pam' /etc/audit/rules.d/*.rules
```

Example output:

```
/etc/audit/rules.d/50-identity.rules:-a always,exit -F arch=b64 -F
path=/etc/pam.conf -F perm=wa -k identity
/etc/audit/rules.d/50-identity.rules:-a always,exit -F arch=b64 -F
dir=/etc/pam.d -F perm=wa -k identity
```

Remediation:

Note:

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor events that modify user/group information.

Example:

```
# printf "
-a always,exit -F arch=b64 -F path=/etc/pam.conf -F perm=wa -k identity
-a always,exit -F arch=b64 -F dir=/etc/pam.d -F perm=wa -k identity
" >> /etc/audit/rules.d/50-identity.rules
```

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000313
2. https://manpages.debian.org/bookworm/manpages/nsswitch.conf.5.en.html
3. https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/6/html/managing_smart_cards/pam_configuration_files
4. NIST SP 800-53 :: AU-12 a
5. NIST SP 800-53A :: AU-12.1 (ii)

Internal Only - General

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	4.8 Log and Alert on Changes to Administrative Group Membership Configure systems to issue a log entry and alert when an account is added to or removed from any group assigned administrative privileges.		●	●

Internal Only - General

6.2.3.18 Ensure discretionary access control permission modification events chmod,fchmod,fchmodat,fchmodat2 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 following commands and system calls effect the permissions, ownership and various attributes of files.

- `chmod`: changes the mode of the file
- `fchmod`: changes the mode of the file referred to by the open file descriptor
- `fchmodat`: change permissions of a file
- `fchmodat2`: change permissions of a file relative to a directory file descriptor

In all cases, an audit record will only be written for non-system user ids and will ignore Daemon events. All audit records will be tagged with the identifier "perm_mod."

Rationale:

Monitoring for changes in file attributes could alert a system administrator to activity that could indicate intruder activity or policy violation.

Satisfies: SRG-OS-000062-GPOS-00031, SRG-OS-000037-GPOS-00015, SRG-OS-000042-GPOS-00020, SRG-OS-000392-GPOS-00172, SRG-OS-000462-GPOS-00206, SRG-OS-000471-GPOS-00215, SRG-OS-000064-GPOS-00033, SRG-OS-000466-GPOS-00210

Audit:

Note:

- Output showing all audited syscalls, e.g. (`-a always,exit -F arch=b64 -S chmod,fchmod,fchmodat,fchmodat2,setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr -F auid>=1000 -F auid!=unset -F key=perm_mod`) is also acceptable. These have been separated by function on the displayed output for clarity.
- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

Internal Only - General

1. Run the following command to verify a rule exists in the running configuration

```
# auditctl -l | grep -Ps -- 'chmod'
```

Verify the output includes:

```
-a always,exit -F arch=b64 -S chmod,fchmod,fchmodat,fchmodat2 -F auid>=1000 -
F auid!=unset
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- 'chmod' /etc/audit/rules.d/*.rules
```

Example output:

```
/etc/audit/rules.d/50-perm_mod.rules:-a always,exit -F arch=b32 -S
chmod,fchmod,fchmodat,fchmodat2 -F auid>=1000 -F auid!=unset -k perm_mod
/etc/audit/rules.d/50-perm_mod.rules:-a always,exit -F arch=b64 -S
chmod,fchmod,fchmodat,fchmodat2 -F auid>=1000 -F auid!=unset -k perm_mod
```

Remediation:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor discretionary access control permission modification events.

Example:

```
# printf "
-a always,exit -F arch=b32 -S chmod,fchmod,fchmodat,fchmodat2 -F auid>=1000 -
F auid!=unset -k perm_mod
-a always,exit -F arch=b64 -S chmod,fchmod,fchmodat,fchmodat2 -F auid>=1000 -
F auid!=unset -k perm_mod
" >> /etc/audit/rules.d/50-perm_mod.rules
```

Note: The `auid>=` value should be set to the same value as `UID_MIN` value in `/etc/login.defs`.

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

Internal Only - General

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000305
2. NIST SP 800-53 :: AU-12 a
3. NIST SP 800-53A :: AU-12.1 (ii)
4. STIG ID: RHEL-08-030490 | RULE ID: SV-230456r810462 | CAT II
5. STIG ID: RHEL-08-030200 | RULE ID: SV-230413r810463 | CAT II
6. STIG ID: RHEL-09-654015 | RULE ID: SV-258177r1045316 | CAT II
7. STIG ID: RHEL-09-654025 | RULE ID: SV-258179r1045322 | CAT II
8. STIG ID: ALMA-09-048640 | RULE ID: SV-269480r1050363 | CAT II
9. STIG ID: ALMA-09-051390 | RULE ID: SV-269505r1050388 | CAT II
10. STIG ID: UBTU-22-654160 | RULE ID: SV-260634r958446 | CAT II
11. STIG ID: UBTU-22-654180 | RULE ID: SV-260638r958446 | CAT II
12. STIG ID: UBTU-24-900140 | RULE ID: SV-270785r1068373 | CAT II
13. STIG ID: UBTU-24-900130 | RULE ID: SV-270784r1068371 | CAT II

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	5.5 Implement Automated Configuration Monitoring Systems 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.		●	●

Internal Only - General

6.2.3.19 Ensure discretionary access control permission modification events chown,fchown,lchown,fchownat 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 following commands and system calls effect the permissions, ownership and various attributes of files.

- `chown`: changes the user and/or group ownership of each given file
- `fchown`: changes the ownership of the file referred to by the open file descriptor
- `fchownat`: change ownership of a file relative to a directory file descriptor
- `lchown`: change the owner and group of a file

In all cases, an audit record will only be written for non-system user ids and will ignore Daemon events. All audit records will be tagged with the identifier "perm_mod."

Rationale:

Monitoring for changes in file attributes could alert a system administrator to activity that could indicate intruder activity or policy violation.

Satisfies: SRG-OS-000062-GPOS-00031, SRG-OS-000037-GPOS-00015, SRG-OS-000042-GPOS-00020, SRG-OS-000392-GPOS-00172, SRG-OS-000462-GPOS-00206, SRG-OS-000471-GPOS-00215, SRG-OS-000064-GPOS-00033, SRG-OS-000466-GPOS-00210

Audit:

Note:

- Output showing all audited syscalls, e.g. (`-a always,exit -F arch=b64 -S chown,fchown,fchownat,lchown,setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr -F auid>=1000 -F auid!=unset -F key=perm_mod`) is also acceptable. These have been separated by function on the displayed output for clarity.
- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

Internal Only - General

1. Run the following command to verify a rule exists in the running configuration

```
# auditctl -l | grep -Ps -- 'chown'
```

Verify the output includes:

```
-a always,exit -F arch=b64 -S chown,fchown,lchown,fchownat -F auid>=1000 -F
auid!=unset
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- 'chown' /etc/audit/rules.d/*.rules
```

Example output:

```
/etc/audit/rules.d/50-perm_mod.rules:-a always,exit -F arch=b32 -S
chown,fchown,lchown,fchownat -F auid>=1000 -F auid!=unset -k perm_mod
/etc/audit/rules.d/50-perm_mod.rules:-a always,exit -F arch=b64 -S
chown,fchown,lchown,fchownat -F auid>=1000 -F auid!=unset -k perm_mod
```

Remediation:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor discretionary access control permission modification events.

Example:

```
# printf "
-a always,exit -F arch=b32 -S chown,fchown,lchown,fchownat -F auid>=1000 -F
auid!=unset -k perm_mod
-a always,exit -F arch=b64 -S chown,fchown,lchown,fchownat -F auid>=1000 -F
auid!=unset -k perm_mod
" >> /etc/audit/rules.d/50-perm_mod.rules
```

Note: The `auid>=` value should be set to the same value as `UID_MIN` value in `/etc/login.defs`.

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

Internal Only - General

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000306
2. NIST SP 800-53 :: AU-12 a
3. NIST SP 800-53A :: AU-12.1 (ii)
4. STIG ID: RHEL-08-030490 | RULE ID: SV-230456r810462 | CAT II
5. STIG ID: RHEL-08-030200 | RULE ID: SV-230413r810463 | CAT II
6. STIG ID: RHEL-09-654015 | RULE ID: SV-258177r1045316 | CAT II
7. STIG ID: RHEL-09-654025 | RULE ID: SV-258179r1045322 | CAT II
8. STIG ID: ALMA-09-048640 | RULE ID: SV-269480r1050363 | CAT II
9. STIG ID: ALMA-09-051390 | RULE ID: SV-269505r1050388 | CAT II
10. STIG ID: UBTU-22-654160 | RULE ID: SV-260634r958446 | CAT II
11. STIG ID: UBTU-22-654180 | RULE ID: SV-260638r958446 | CAT II
12. STIG ID: UBTU-24-900140 | RULE ID: SV-270785r1068373 | CAT II
13. STIG ID: UBTU-24-900130 | RULE ID: SV-270784r1068371 | CAT II

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 Collect Detailed Audit Logs Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.		●	●
v7	5.5 Implement Automated Configuration Monitoring Systems 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.		●	●

Internal Only - General

Internal Only - General

6.2.3.20 Ensure discretionary access control permission modification events setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr 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 following commands and system calls effect the permissions, ownership and various attributes of files.

- `setxattr`
- `lsetxattr`
- `fsetxattr`
- `removexattr`
- `lremovexattr`
- `fremovexattr`

In all cases, an audit record will only be written for non-system user ids and will ignore Daemon events. All audit records will be tagged with the identifier "perm_mod."

Rationale:

Monitoring for changes in file attributes could alert a system administrator to activity that could indicate intruder activity or policy violation.

Satisfies: SRG-OS-000062-GPOS-00031, SRG-OS-000037-GPOS-00015, SRG-OS-000042-GPOS-00020, SRG-OS-000392-GPOS-00172, SRG-OS-000462-GPOS-00206, SRG-OS-000471-GPOS-00215, SRG-OS-000064-GPOS-00033, SRG-OS-000466-GPOS-00210

Audit:

Note:

- Output showing all audited syscalls, e.g. (`-a always,exit -F arch=b64 -S chmod,fchmod,fchmodat,fchmodat2,setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr -F auid>=1000 -F auid!=unset -F key=perm_mod`) is also acceptable. These have been separated by function on the displayed output for clarity.

Internal Only - General

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Run the following command to verify a rule exists in the running configuration

```
# auditctl -l | grep -Ps -- 'xattr'
```

Verify the output includes:

```
-a always,exit -F arch=b64 -S
setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr -F
auid>=1000 -F auid!=unset
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- 'xattr' /etc/audit/rules.d/*.rules
```

Example output:

```
/etc/audit/rules.d/50-perm_mod.rules:-a always,exit -F arch=b32 -S
setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr -F
auid>=1000 -F auid!=unset -k perm_mod
/etc/audit/rules.d/50-perm_mod.rules:-a always,exit -F arch=b64 -S
setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr -F
auid>=1000 -F auid!=unset -k perm_mod
```

Remediation:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor discretionary access control permission modification events.

Example:

```
# printf "
-a always,exit -F arch=b32 -S
setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr -F
auid>=1000 -F auid!=unset -k perm_mod
-a always,exit -F arch=b64 -S
setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr -F
auid>=1000 -F auid!=unset -k perm_mod
" >> /etc/audit/rules.d/50-perm_mod.rules
```

Note: The `auid>=` value should be set to the same value as `UID_MIN` value in `/etc/login.defs`.

2. Merge and load the rules into active configuration:

Internal Only - General

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000307
2. NIST SP 800-53 :: AU-12 a
3. NIST SP 800-53A :: AU-12.1 (ii)
4. STIG ID: RHEL-08-030490 | RULE ID: SV-230456r810462 | CAT II
5. STIG ID: RHEL-08-030200 | RULE ID: SV-230413r810463 | CAT II
6. STIG ID: RHEL-09-654015 | RULE ID: SV-258177r1045316 | CAT II
7. STIG ID: RHEL-09-654025 | RULE ID: SV-258179r1045322 | CAT II
8. STIG ID: ALMA-09-048640 | RULE ID: SV-269480r1050363 | CAT II
9. STIG ID: ALMA-09-051390 | RULE ID: SV-269505r1050388 | CAT II
10. STIG ID: UBTU-22-654160 | RULE ID: SV-260634r958446 | CAT II
11. STIG ID: UBTU-22-654180 | RULE ID: SV-260638r958446 | CAT II
12. STIG ID: UBTU-24-900140 | RULE ID: SV-270785r1068373 | CAT II
13. STIG ID: UBTU-24-900130 | RULE ID: SV-270784r1068371 | CAT II

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 Collect Detailed Audit Logs Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.		●	●

Internal Only - General

Controls Version	Control	IG 1	IG 2	IG 3
v7	5.5 Implement Automated Configuration Monitoring Systems 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.		●	●

Internal Only - General

6.2.3.21 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.

Satisfies: SRG-OS-000062-GPOS-00031, SRG-OS-000037-GPOS-00015, SRG-OS-000042-GPOS-00020, SRG-OS-000062-GPOS-00031, SRG-OS-000392-GPOS-00172, SRG-OS-000462-GPOS-00206, SRG-OS-000471-GPOS-00215

Audit:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps -- '\-S mount'
```

Verify the output includes:

```
-a always,exit -F arch=b64 -S mount -F auid>=1000 -F auid!=unset
```

2. Run the following command to verify a rule exists in a rules file:

Internal Only - General

```
# grep -Ps -- '\-S mount' /etc/audit/rules.d/*.rules
```

Example output:

```
/etc/audit/rules.d/50-mounts.rules:-a always,exit -F arch=b32 -S mount -F
auid>=1000 -F auid!=unset -k mounts
/etc/audit/rules.d/50-mounts.rules:-a always,exit -F arch=b64 -S mount -F
auid>=1000 -F auid!=unset -k mounts
```

Remediation:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor successful file system mounts.

Example:

```
# printf "
-a always,exit -F arch=b32 -S mount -F auid>=1000 -F auid!=unset -k mounts
-a always,exit -F arch=b64 -S mount -F auid>=1000 -F auid!=unset -k mounts
" >> /etc/audit/rules.d/50-mounts.rules
```

Note: The `auid>=` value should be set to the same value as `UID_MIN` value in `/etc/login.defs`.

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000337
2. NIST SP 800-53 :: AU-12 a
3. NIST SP 800-53A :: AU-12.1 (ii)
4. STIG ID: RHEL-08-030302 | RULE ID: SV-230425r627750 | CAT II

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

Internal Only - General

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 Collect Detailed Audit Logs Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.		●	●
v7	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.		●	●

Internal Only - General

6.2.3.22 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.

- `/var/run/utmp` - tracks all currently logged in users.
- `/var/log/wtmp` - file tracks logins, logouts, shutdown, and reboot events.
- `/var/log/btmp` - keeps track of failed login attempts and can be read by entering the command `/usr/bin/last -f /var/log/btmp`.

All audit records will be tagged with the identifier "session."

Rationale:

Monitoring these files for changes could alert a system administrator to logins occurring at unusual hours, which could indicate intruder activity (i.e. a user logging in at a time when they do not normally log in).

Audit:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps --
'(\/var\/run\/utmp|\/var\/log\/wtmp|\/var\/log\/btmp)'
```

Verify the output matches:

```
-a always,exit -F path=/var/run/utmp -F perm=wa
-a always,exit -F path=/var/log/wtmp -F perm=wa
-a always,exit -F path=/var/log/btmp -F perm=wa
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- '(\/var\/run\/utmp|\/var\/log\/wtmp|\/var\/log\/btmp)'
/etc/audit/rules.d/*.rules
```

Verify the output matches:

Internal Only - General

```
/etc/audit/rules.d/50-session.rules:-a always,exit -F arch=b64 -F
path=/var/run/utmp -F perm=wa -k session
/etc/audit/rules.d/50-session.rules:-a always,exit -F arch=b64 -F
path=/var/log/wtmp -F perm=wa -k session
/etc/audit/rules.d/50-session.rules:-a always,exit -F arch=b64 -F
path=/var/log/btmp -F perm=wa -k session
```

Remediation:

Note:

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor session initiation information.

Example:

```
# printf "
-a always,exit -F arch=b64 -F path=/var/run/utmp -F perm=wa -k session
-a always,exit -F arch=b64 -F path=/var/log/wtmp -F perm=wa -k session
-a always,exit -F arch=b64 -F path=/var/log/btmp -F perm=wa -k session
" >> /etc/audit/rules.d/50-session.rules
```

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000332
2. NIST SP 800-53 Rev. 5: AU-3, AU-3
3. STIG ID: UBTU-24-900600 | RULE ID: SV-270811r1066922 | CAT II
4. STIG ID: UBTU-22-654205 | RULE ID: SV-260643r991581 | CAT II
5. STIG ID: UBTU-24-900590 | RULE ID: SV-270810r1066919 | CAT II
6. STIG ID: UBTU-22-654200 | RULE ID: SV-260642r991581 | CAT II
7. STIG ID: UBTU-24-900610 | RULE ID: SV-270812r1066925 | CAT II
8. STIG ID: UBTU-22-654195 | RULE ID: SV-260641r991581 | CAT II

Internal Only - General

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	4.9 Log and Alert on Unsuccessful Administrative Account Login Configure systems to issue a log entry and alert on unsuccessful logins to an administrative account.		●	●
v7	16.13 Alert on Account Login Behavior Deviation Alert when users deviate from normal login behavior, such as time-of-day, workstation location and duration.			●

Internal Only - General

6.2.3.23 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.

- `/var/log/lastlog` - maintain records of the last time a user successfully logged in.
- `/var/run/faillock` - directory maintains records of login failures via the `pam_faillock` module.

Rationale:

Monitoring login/logout events could provide a system administrator with information associated with brute force attacks against user logins.

Satisfies: SRG-OS-000062-GPOS-00031, SRG-OS-000037-GPOS-00015, SRG-OS-000042-GPOS-00020, SRG-OS-000062-GPOS-00031, SRG-OS-000392-GPOS-00172, SRG-OS-000462-GPOS-00206, SRG-OS-000471-GPOS-00215, SRG-OS-000473-GPOS-00218

Audit:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps -- '(lastlog|faillock)'
```

Verify the output includes:

```
-a always,exit -F path=/var/log/lastlog -F perm=wa
-a always,exit -F path=/var/run/faillock -F perm=wa
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- '(lastlog|faillock)' /etc/audit/rules.d/*.rules
```

Example output:

Internal Only - General

```
/etc/audit/rules.d/50-login.rules:-a always,exit -F arch=b64 -F
path=/var/log/lastlog -F perm=wa -k logins
/etc/audit/rules.d/50-login.rules:-a always,exit -F arch=b64 -F
path=/var/run/faillock -F perm=wa -k logins
```

Remediation:

Note:

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor login and logout events.

Example:

```
# printf "
-a always,exit -F arch=b64 -F path=/var/log/lastlog -F perm=wa -k logins
-a always,exit -F arch=b64 -F path=/var/run/faillock -F perm=wa -k logins
" >> /etc/audit/rules.d/50-login.rules
```

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000329
2. NIST SP 800-53 :: AU-12 a
3. NIST SP 800-53A :: AU-12.1 (ii)
4. STIG ID: RHEL-08-030600 | RULE ID: SV-230467r627750 | CAT II
5. STIG ID: RHEL-09-654255 | RULE ID: SV-258225r1014990 | CAT II
6. STIG ID: ALMA-09-047540 | RULE ID: SV-269470r1050353 | CAT II
7. STIG ID: UBTU-24-900260 | RULE ID: SV-270797r1066880 | CAT II
8. STIG ID: UBTU-22-654215 | RULE ID: SV-260645r958446 | CAT II

Internal Only - General

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	4.9 Log and Alert on Unsuccessful Administrative Account Login Configure systems to issue a log entry and alert on unsuccessful logins to an administrative account.		●	●
v7	16.11 Lock Workstation Sessions After Inactivity Automatically lock workstation sessions after a standard period of inactivity.	●	●	●
v7	16.13 Alert on Account Login Behavior Deviation Alert when users deviate from normal login behavior, such as time-of-day, workstation location and duration.			●

Internal Only - General

6.2.3.24 Ensure unlink 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:

- `unlink` - remove a file
- `unlinkat` - remove a file attribute

  system calls and tags them with the identifier "delete".

Rationale:

Monitoring these calls from non-privileged users could provide a system administrator with evidence that inappropriate removal of files and file attributes associated with protected files is occurring. While this audit option will look at all events, system administrators will want to look for specific privileged files that are being deleted or altered.

Satisfies: SRG-OS-000062-GPOS-00031, SRG-OS-000037-GPOS-00015, SRG-OS-000042-GPOS-00020, SRG-OS-000392-GPOS-00172, SRG-OS-000462-GPOS-00206, SRG-OS-000471-GPOS-00215

Audit:

Note:

- Output showing all audited syscalls, e.g. (`-a always,exit -S unlink,unlinkat,rename,renameat,renameat2 -F auid>=1000 -F auid!=unset`) is also acceptable. These have been separated by function on the displayed output for clarity.
- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps -- 'unlink'
```

Verify the output includes:

Internal Only - General

```
-a always,exit -S unlink,unlinkat -F auid>=1000 -F auid!=unset
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- 'unlink' /etc/audit/rules.d/*.rules
```

Example output:

```
/etc/audit/rules.d/50-delete.rules:-a always,exit -F arch=b32 -S
unlink,unlinkat -F auid>=1000 -F auid!=unset -k delete
/etc/audit/rules.d/50-delete.rules:-a always,exit -F arch=b64 -S
unlink,unlinkat -F auid>=1000 -F auid!=unset -k delete
```

Remediation:

> **Note:** The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure. For best performance, the arch field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the arch will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

Edit or create a file in the `/etc/audit/rules.d/` directory, ending in .rules extension, with the relevant rules to monitor file deletion events by users.

Example:

```
# printf "
-a always,exit -F arch=b32 -S unlink,unlinkat -F auid>=1000 -F auid!=unset -k
delete
-a always,exit -F arch=b64 -S unlink,unlinkat -F auid>=1000 -F auid!=unset -k
delete
" >> /etc/audit/rules.d/50-delete.rules
```

> **Note:** The `auid>=` value should be set to the same value as `UID_MIN` value in `/etc/login.defs`.

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000344
2. NIST SP 800-53 :: AU-12 a
3. NIST SP 800-53A :: AU-12.1 (ii)
4. STIG ID: RHEL-08-030361 | RULE ID: SV-230439r810465 | CAT II
5. STIG ID: RHEL-09-654065 | RULE ID: SV-258187r1045346 | CAT II

Internal Only - General

6. STIG ID: ALMA-09-048970 | RULE ID: SV-269483r1050366 | CAT II
7. STIG ID: UBTU-22-654185 | RULE ID: SV-260639r991577 | CAT II
8. STIG ID: UBTU-24-900540 | RULE ID: SV-270809r1068388 | CAT II

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.	●	●	●

Internal Only - General

6.2.3.25 Ensure rename 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:

- `rename`: rename a file
- `renameat`: rename a file attribute
- `renameat2`: rename a file relative to directory file descriptors

system calls and tags them with the identifier "delete".

Rationale:

Monitoring these calls from non-privileged users could provide a system administrator with evidence that inappropriate removal of files and file attributes associated with protected files is occurring. While this audit option will look at all events, system administrators will want to look for specific privileged files that are being deleted or altered.

Satisfies: SRG-OS-000062-GPOS-00031, SRG-OS-000037-GPOS-00015, SRG-OS-000042-GPOS-00020, SRG-OS-000392-GPOS-00172, SRG-OS-000462-GPOS-00206, SRG-OS-000471-GPOS-00215

Audit:

Note:

- Output showing all audited syscalls, e.g. (`-a always,exit -S unlink,unlinkat,rename,renameat,renameat2 -F auid>=1000 -F auid!=unset`) is also acceptable. These have been separated by function on the displayed output for clarity.
- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps -- 'rename'
```

Verify the output includes:

Internal Only - General

```
-a always,exit -S rename,renameat,renameat2 -F auid>=1000 -F auid!=unset
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- 'rename' /etc/audit/rules.d/*.rules
```

Example output:

```
/etc/audit/rules.d/50-delete.rules:-a always,exit -F arch=b32 -S
rename,renameat,renameat2 -F auid>=1000 -F auid!=unset -k delete
/etc/audit/rules.d/50-delete.rules:-a always,exit -F arch=b64 -S
rename,renameat,renameat2 -F auid>=1000 -F auid!=unset -k delete
```

Remediation:

Note:

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor file deletion events by users.

Example:

```
# printf "
-a always,exit -F arch=b32 -S rename,renameat,renameat2 -F auid>=1000 -F
auid!=unset -k delete
-a always,exit -F arch=b64 -S rename,renameat,renameat2 -F auid>=1000 -F
auid!=unset -k delete
" >> /etc/audit/rules.d/50-delete.rules
```

Note: The `auid>=` value should be set to the same value as `UID_MIN` value in `/etc/login.defs`.

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000331

Internal Only - General

2. NIST SP 800-53 :: AU-12 a
3. NIST SP 800-53A :: AU-12.1 (ii)
4. STIG ID: RHEL-08-030361 | RULE ID: SV-230439r810465 | CAT II
5. STIG ID: RHEL-09-654065 | RULE ID: SV-258187r1045346 | CAT II
6. STIG ID: ALMA-09-048970 | RULE ID: SV-269483r1050366 | CAT II
7. STIG ID: UBTU-22-654185 | RULE ID: SV-260639r991577 | CAT II
8. STIG ID: UBTU-24-900540 | RULE ID: SV-270809r1068388 | CAT II

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.	●	●	●

Internal Only - General

6.2.3.26 Ensure events that modify the system's Mandatory Access Controls are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Monitor AppArmor, an implementation of mandatory access controls. The parameters below monitor any write access (potential additional, deletion or modification of files in the directory) or attribute changes to the `/etc/apparmor/` and `/etc/apparmor.d/` directories.

Note: If a different Mandatory Access Control method is used, changes to the corresponding directories should be audited.

Rationale:

Changes to files in the `/etc/apparmor/` and `/etc/apparmor.d/` 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:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps -- 'apparmor'
```

Verify the output includes:

```
-a always,exit -F dir=/etc/apparmor -F perm=wa
-a always,exit -F dir=/etc/apparmor.d -F perm=wa
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- 'apparmor' /etc/audit/rules.d/*.rules
```

Example output:

Internal Only - General

```
/etc/audit/rules.d/50-MAC-policy.rules:-a always,exit -F arch=b64 -F
dir=/etc/apparmor -F perm=wa -k MAC-policy
/etc/audit/rules.d/50-MAC-policy.rules:-a always,exit -F arch=b64 -F
dir=/etc/apparmor.d -F perm=wa -k MAC-policy
```

Remediation:

Note:

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor events that modify the system's Mandatory Access Controls.

Example:

```
# printf "
-a always,exit -F arch=b64 -F dir=/etc/apparmor -F perm=wa -k MAC-policy
-a always,exit -F arch=b64 -F dir=/etc/apparmor.d -F perm=wa -k MAC-policy
" >> /etc/audit/rules.d/50-MAC-policy.rules
```

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000321
2. NIST SP 800-53 Rev. 5: AU-3, CM-6

Internal Only - General

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	5.5 Implement Automated Configuration Monitoring Systems 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.		●	●

Internal Only - General

6.2.3.27 Ensure successful and unsuccessful attempts to use the chcon command are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

The operating system must generate audit records for successful/unsuccessful uses of the `chcon` command.

Rationale:

The `chcon` command is used to change file security context. Without generating audit records that are specific to the security and needs of the organization, it would be difficult to establish, correlate, and investigate the events relating to an incident or identify those responsible for one.

Audit records can be generated from various components within the information system (e.g., module or policy filter).

Satisfies: SRG-OS-000062-GPOS-00031, SRG-OS-000037-GPOS-00015, SRG-OS-000042-GPOS-00020, SRG-OS-000062-GPOS-00031, SRG-OS-000392-GPOS-00172, SRG-OS-000462-GPOS-00206, SRG-OS-000468-GPOS-00212, SRG-OS-000471-GPOS-00215

Audit:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Run the following command to verify a rule exits in the running configuration:

```
# auditctl -l | grep -Ps -- '^\h*[^#\n\r]+\h*\/usr\/bin\/chcon'
```

Verify the output includes:

```
-a always,exit -F path=/usr/bin/chcon -F perm=x -F auid>=1000 -F auid!=unset
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- '^\h*[^#\n\r]+\h*\/usr\/bin\/chcon' /etc/audit/rules.d/*.rules
```

Example output:

Internal Only - General

```
/etc/audit/rules.d/50-perm_chng.rules:-a always,exit -F arch=b64 -F
path=/usr/bin/chcon -F perm=x -F auid>=1000 -F auid!=unset -k perm_chng
```

Remediation:

Note:

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor successful and unsuccessful attempts to use the `chcon` command.

Example:

```
# printf "
-a always,exit -F arch=b64 -F path=/usr/bin/chcon -F perm=x -F auid>=1000 -F
auid!=unset -k perm_chng
" >> /etc/audit/rules.d/50-perm_chng.rules
```

Note: The `auid>=` value should be set to the same value as `UID_MIN` value in `/etc/login.defs`.

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000334
2. NIST SP 800-53 :: AU-12 a
3. NIST SP 800-53A :: AU-12.1 (ii)
4. STIG ID: RHEL-08-030260 | RULE ID: SV-230419r627750 | CAT II
5. STIG ID: RHEL-09-654045 | RULE ID: SV-258183r1045334 | CAT II
6. STIG ID: ALMA-09-048420 | RULE ID: SV-269478r1050361 | CAT II
7. STIG ID: UBTU-22-654025 | RULE ID: SV-260607r958446 | CAT II
8. STIG ID: UBTU-24-900210 | RULE ID: SV-270792r1066865 | CAT II

Internal Only - General

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

System call structure

For performance (`man 7 audit.rules`) reasons it is preferable to have all the system calls on one line. However, your configuration may have them on one line each or some other combination. This is important to understand for both the auditing and remediation sections as the examples given are optimized for performance as per the man page.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 Collect Audit Logs Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.	●	●	●
v7	6.2 Activate audit logging Ensure that local logging has been enabled on all systems and networking devices.	●	●	●

Internal Only - General

6.2.3.28 Ensure successful and unsuccessful attempts to use the setfacl command are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

The operating system must generate audit records for successful/unsuccessful uses of the `setfacl` command

Rationale:

This utility sets Access Control Lists (ACLs) of files and directories. Without generating audit records that are specific to the security and mission needs of the organization, it would be difficult to establish, correlate, and investigate the events relating to an incident or identify those responsible for one.

Audit records can be generated from various components within the information system (e.g., module or policy filter).

Audit:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps '^\h*[^#\n\r]+\h*\/usr\/bin\/setfacl'
```

Verify the output includes:

```
-a always,exit -F path=/usr/bin/setfacl -F perm=x -F auid>=1000 -F
auid!=unset
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps '^\h*[^#\n\r]+\h*\/usr\/bin\/setfacl' /etc/audit/rules.d/*.rules
```

Example output:

```
/etc/audit/rules.d/50-perm_chng.rules:-a always,exit -F arch=b64 -F
path=/usr/bin/setfacl -F perm=x -F auid>=1000 -F auid!=unset -k perm_chng
```

Remediation:

Note:

Internal Only - General

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor successful and unsuccessful attempts to use the `setfacl` command.

Example:

```
# printf "
-a always,exit -F arch=b64 -F path=/usr/bin/setfacl -F perm=x -F auid>=1000 -
F auid!=unset -k perm_chng
" >> /etc/audit/rules.d/50-perm_chng.rules
```

Note: The `auid>=` value should be set to the same value as `UID_MIN` value in `/etc/login.defs`.

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000335
2. NIST SP 800-53 Rev. 5: AU-2, AU-12, SI-5
3. STIG ID: RHEL-08-030330 | RULE ID: SV-230435r627750 | CAT II
4. STIG ID: RHEL-09-654040 | RULE ID: SV-258182r1045331 | CAT II
5. STIG ID: ALMA-09-050180 | RULE ID: SV-269494r1050377 | CAT II
6. STIG ID: UBTU-22-654085 | RULE ID: SV-260619r958446 | CAT II
7. STIG ID: UBTU-24-900230 | RULE ID: SV-270794r1066871 | CAT II

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.	●	●	●

Internal Only - General

Controls Version	Control	IG 1	IG 2	IG 3
v7	6.2 Activate audit logging Ensure that local logging has been enabled on all systems and networking devices.	●	●	●

Internal Only - General

6.2.3.29 Ensure successful and unsuccessful attempts to use the chacl command are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

The operating system must generate audit records for successful/unsuccessful uses of the `chacl` command.

`chacl` is an IRIX-compatibility command, and is maintained for those users who are familiar with its use from either XFS or IRIX.

Rationale:

`chacl` changes the ACL(s) for a file or directory. Without generating audit records that are specific to the security and mission needs of the organization, it would be difficult to establish, correlate, and investigate the events relating to an incident or identify those responsible for one.

Audit records can be generated from various components within the information system (e.g., module or policy filter).

Satisfies: SRG-OS-000062-GPOS-00031, SRG-OS-000037-GPOS-00015, SRG-OS-000042-GPOS-00020, SRG-OS-000062-GPOS-00031, SRG-OS-000392-GPOS-00172, SRG-OS-000462-GPOS-00206, SRG-OS-000471-GPOS-00215, SRG-OS-000466-GPOS-00210

Audit:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps -- '^\h*[^#\n\r]+\h*\/usr\/bin\/chacl'
```

Verify the output includes:

```
-a always,exit -F path=/usr/bin/chacl -F perm=x -F auid>=1000 -F auid!=unset
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- '^\h*[^#\n\r]+\h*\/usr\/bin\/chacl' /etc/audit/rules.d/*.rules
```

Example output:

Internal Only - General

```
/etc/audit/rules.d/50-perm_chng.rules:-a always,exit -F arch=b64 -F
path=/usr/bin/chacl -F perm=x -F auid>=1000 -F auid!=unset -k perm_chng
```

Remediation:

Note:

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor successful and unsuccessful attempts to use the `chacl` command.

Example:

```
# printf "
-a always,exit -F arch=b64 -F path=/usr/bin/chacl -F perm=x -F auid>=1000 -F
auid!=unset -k perm_chng
" >> /etc/audit/rules.d/50-perm_chng.rules
```

Note: The `auid>=` value should be set to the same value as `UID_MIN` value in `/etc/login.defs`.

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000333
2. NIST SP 800-53 Rev. 5: AU-2, AU-12, SI-5
3. NIST SP 800-53A :: AU-12.1 (ii)
4. STIG ID: RHEL-08-030570 | RULE ID: SV-230464r627750 | CAT II
5. STIG ID: RHEL-09-654035 | RULE ID: SV-258181r1045328 | CAT II
6. STIG ID: ALMA-09-048200 | RULE ID: SV-269476r1050359 | CAT II
7. STIG ID: UBTU-22-654015 | RULE ID: SV-260605r958446 | CAT II
8. STIG ID: UBTU-24-900240 | RULE ID: SV-270795r1066874 | CAT II

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 Collect Audit Logs Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.	●	●	●
v7	6.2 Activate audit logging Ensure that local logging has been enabled on all systems and networking devices.	●	●	●

Internal Only - General

6.2.3.30 Ensure successful and unsuccessful attempts to use the usermod command are collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

The operating system must generate audit records for successful/unsuccessful uses of the `usermod` command.

Rationale:

The `usermod` command modifies the system account files to reflect the changes that are specified on the command line. Without generating audit records that are specific to the security and needs of the organization, it would be difficult to establish, correlate, and investigate the events relating to an incident or identify those responsible for one.

Audit records can be generated from various components within the information system (e.g., module or policy filter).

Satisfies: SRG-OS-000062-GPOS-00031, SRG-OS-000037-GPOS-00015, SRG-OS-000042-GPOS-00020, SRG-OS-000062-GPOS-00031, SRG-OS-000392-GPOS-00172, SRG-OS-000462-GPOS-00206, SRG-OS-000471-GPOS-00215, SRG-OS-000466-GPOS-00210

Audit:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps -- '^\h*[^#\n\r]+\h*\/usr\/sbin\/usermod'
```

Verify the output includes:

```
-a always,exit -F path=/usr/sbin/usermod -F perm=x -F auid>=1000 -F
auid!=unset
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- '^\h*[^#\n\r]+\h*\/usr\/sbin\/usermod'
/etc/audit/rules.d/*.rules
```

Example output:

Internal Only - General

```
/etc/audit/rules.d/50-usermod.rules:-a always,exit -F arch=b64 -F
path=/usr/sbin/usermod -F perm=x -F auid>=1000 -F auid!=unset -k usermod
```

Remediation:

Note:

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor successful and unsuccessful attempts to use the `usermod` command.

Example:

```
# printf "
-a always,exit -F arch=b64 -F path=/usr/sbin/usermod -F perm=x -F auid>=1000
-F auid!=unset -k usermod
" >> /etc/audit/rules.d/50-usermod.rules
```

Note: The `auid>=` value should be set to the same value as `UID_MIN` value in `/etc/login.defs`.

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000336
2. NIST SP 800-53 :: AU-12 a
3. NIST SP 800-53A :: AU-12.1 (ii)
4. STIG ID: RHEL-08-030560 | RULE ID: SV-230463r627750 | CAT II
5. STIG ID: RHEL-09-654175 | RULE ID: SV-258209r1045412 | CAT II
6. STIG ID: ALMA-09-051280 | RULE ID: SV-269504r1050387 | CAT II
7. STIG ID: UBTU-22-654125 | RULE ID: SV-260627r958446 | CAT II
8. STIG ID: UBTU-24-900310 | RULE ID: SV-270802r1066895 | CAT II

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 Collect Audit Logs Collect audit logs. Ensure that logging, per the enterprise's audit log management process, has been enabled across enterprise assets.	●	●	●
v7	6.2 Activate audit logging Ensure that local logging has been enabled on all systems and networking devices.	●	●	●

Internal Only - General

6.2.3.31 Ensure kernel module loading unloading and modification is collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Monitor the loading and unloading of kernel modules. All the loading / listing / dependency checking of modules is done by `kmod` via symbolic links.

Rationale:

Monitoring the use of all the various ways to manipulate kernel modules could provide system administrators with evidence that an unauthorized change was made to a kernel module, possibly compromising the security of the system.

Satisfies: SRG-OS-000062-GPOS-00031, SRG-OS-000037-GPOS-00015, SRG-OS-000042-GPOS-00020, SRG-OS-000062-GPOS-00031, SRG-OS-000392-GPOS-00172, SRG-OS-000462-GPOS-00206, SRG-OS-000471-GPOS-00215, SRG-OS-000471-GPOS-00216, SRG-OS-000477-GPOS-00222

Audit:

Note: The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps -- '^\h*[^#\n\r]+\h*\/usr\/bin\/kmod'
```

Verify the output includes:

```
-a always,exit -F path=/usr/bin/kmod -F perm=x -F auid>=1000 -F auid!=unset
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- '^\h*[^#\n\r]+\h*\/usr\/bin\/kmod' /etc/audit/rules.d/*.rules
```

Example output:

```
/etc/audit/rules.d/50-kernel_modules.rules:-a always,exit -F arch=b64 -F
path=/usr/bin/kmod -F perm=x -F auid>=1000 -F auid!=unset -k kernel_modules
```

Symlink audit

Run the following script to audit if the symlinks `kmod` accepts are indeed pointing at it:

Internal Only - General

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

{
   a_files=("/usr/sbin/lsmod" "/usr/sbin/rmmod" "/usr/sbin/insmod"
"/usr/sbin/modinfo" "/usr/sbin/modprobe" "/usr/sbin/depmod")
   for l_file in "${a_files[@]}"; do
      if [ "$(readlink -f "$l_file")" = "$(readlink -f /bin/kmod)" ]; then
         printf "OK: \"$l_file\"\n"
      else
         printf "Issue with symlink for file: \"$l_file\"\n"
      fi
   done
}
```

Verify the output states `OK`. If there is a symlink pointing to a different location it should be investigated

Remediation:

Note:

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor kernel module modification.

Example:

```
# printf "
-a always,exit -F arch=b64 -F path=/usr/bin/kmod -F perm=x -F auid>=1000 -F
auid!=unset -k kernel_modules
" >> /etc/audit/rules.d/50-kernel_modules.rules
```

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000327
2. NIST SP 800-53 :: AU-12 a

Internal Only - General

3. NIST SP 800-53A :: AU-12.1 (ii)

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.	●	●	●

Internal Only - General

6.2.3.32 Ensure kernel "init_module" and "finit_module" loading unloading and modification is collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

The following system calls control loading and unloading of modules:

- `init_module` - load a module
- `finit_module` - load a module (used when the overhead of using cryptographically signed modules to determine the authenticity of a module can be avoided)

Any execution of the loading and unloading module programs and system calls will trigger an audit record with an identifier of `modules`.

Rationale:

Monitoring the use of all the various ways to manipulate kernel modules could provide system administrators with evidence that an unauthorized change was made to a kernel module, possibly compromising the security of the system.

Satisfies: SRG-OS-000062-GPOS-00031, SRG-OS-000037-GPOS-00015, SRG-OS-000042-GPOS-00020, SRG-OS-000062-GPOS-00031, SRG-OS-000392-GPOS-00172, SRG-OS-000462-GPOS-00206, SRG-OS-000471-GPOS-00215, SRG-OS-000471-GPOS-00216, SRG-OS-000477-GPOS-00222

Audit:

Note:

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- Output showing all the audited syscalls combined (e.g. `-a always,exit -F arch=b64 -S init_module,finit_module -F auid>=1000 -F auid!=unset -k kernel_modules`) is also acceptable.

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps -- 'init_module'
```

Verify the output includes both arch variants for `init_module` and `finit_module`:

Internal Only - General

```
-a always,exit -F arch=b32 -S init_module,finit_module -F auid>=1000 -F
auid!=unset -k kernel_modules
-a always,exit -F arch=b64 -S init_module,finit_module -F auid>=1000 -F
auid!=unset -k kernel_modules
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- 'init_module' /etc/audit/rules.d/*.rules
```

Example output:

```
/etc/audit/rules.d/50-kernel_modules.rules:-a always,exit -F arch=b32 -S
init_module,finit_module -F auid>=1000 -F auid!=unset -k kernel_modules
/etc/audit/rules.d/50-kernel_modules.rules:-a always,exit -F arch=b64 -S
init_module,finit_module -F auid>=1000 -F auid!=unset -k kernel_modules
```

Remediation:

Note:

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor kernel module loading.

Example:

```
# printf "
-a always,exit -F arch=b32 -S init_module,finit_module -F auid>=1000 -F
auid!=unset -k kernel_modules
-a always,exit -F arch=b64 -S init_module,finit_module -F auid>=1000 -F
auid!=unset -k kernel_modules
" >> /etc/audit/rules.d/50-kernel_modules.rules
```

Note: The `auid>=` value should be set to the same value as `UID_MIN` value in `/etc/login.defs`.

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

Internal Only - General

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000326
2. NIST SP 800-53 :: AU-12 a
3. NIST SP 800-53A :: AU-12.1 (ii)
4. STIG ID: RHEL-08-030360 | RULE ID: SV-230438r1017241 | CAT II
5. STIG ID: RHEL-09-654080 | RULE ID: SV-258198r1045379 | CAT II
6. STIG ID: ALMA-09-046660 | RULE ID: SV-269466r1050349 | CAT II
7. STIG ID: UBTU-22-654175 | RULE ID: SV-260637r958446 | CAT II
8. STIG ID: UBTU-24-900340 | RULE ID: SV-270805r1068384 | CAT II

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.	●	●	●

Internal Only - General

6.2.3.33 Ensure kernel "delete_module" loading unloading and modification is collected (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Monitor the loading and unloading of kernel modules. All the loading / listing / dependency checking of modules is done by `kmod` via symbolic links.

The following system calls control loading and unloading of modules:

- `delete_module` - delete a module

Any execution of the loading and unloading module programs and system calls will trigger an audit record with an identifier of `modules`.

Rationale:

Monitoring the use of all the various ways to manipulate kernel modules could provide system administrators with evidence that an unauthorized change was made to a kernel module, possibly compromising the security of the system.

Satisfies: SRG-OS-000062-GPOS-00031, SRG-OS-000037-GPOS-00015, SRG-OS-000042-GPOS-00020, SRG-OS-000062-GPOS-00031, SRG-OS-000392-GPOS-00172, SRG-OS-000462-GPOS-00206, SRG-OS-000471-GPOS-00215, SRG-OS-000471-GPOS-00216, SRG-OS-000477-GPOS-00222

Audit:

Note:

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- Output showing all the audited syscalls, e.g. (`-a always,exit -F arch=b64 -S init_module,finit_module,delete_module -F auid>=1000 -F auid!=unset -k kernel_modules`) is also acceptable. These have been separated by function on the displayed output for clarity.

1. Run the following command to verify a rule exists in the running configuration:

```
# auditctl -l | grep -Ps -- 'delete_module'
```

Verify the output includes `delete_module`:

Internal Only - General

```
-a always,exit -S init_module,finit_module,delete_module -F auid>=1000 -F
auid!=unset
```

2. Run the following command to verify a rule exists in a rules file:

```
# grep -Ps -- 'delete_module' /etc/audit/rules.d/*.rules
```

Example output:

```
/etc/audit/rules.d/50-kernel_modules.rules:-a always,exit -F arch=b32 -S
delete_module -F auid>=1000 -F auid!=unset -k kernel_modules
/etc/audit/rules.d/50-kernel_modules.rules:-a always,exit -F arch=b64 -S
delete_module -F auid>=1000 -F auid!=unset -k kernel_modules
```

Remediation:

Note:

- The `-w` option is deprecated and the rule should be updated in accordance with the Remediation Procedure.
- For best performance, the `arch` field should be supplied in the rule. The individual permissions will cause the selection of specific system calls that use that kind of access. Not supplying the `arch` will cause the selection of all system calls which will affect performance as all system calls will be evaluated.

1. Edit or create a file in the `/etc/audit/rules.d/` directory, ending in `.rules` extension, with the relevant rules to monitor kernel module modification.

Example:

```
# printf "
-a always,exit -F arch=b32 -S delete_module -F auid>=1000 -F auid!=unset -k
kernel_modules
-a always,exit -F arch=b64 -S delete_module -F auid>=1000 -F auid!=unset -k
kernel_modules
" >> /etc/audit/rules.d/50-kernel_modules.rules
```

Note: The `auid>=` value should be set to the same value as `UID_MIN` value in `/etc/login.defs`.

2. Merge and load the rules into active configuration:

```
# augenrules --load
```

3. Check if reboot is required:

Internal Only - General

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000325
2. NIST SP 800-53 :: AU-12 a
3. NIST SP 800-53A :: AU-12.1 (ii)
4. STIG ID: RHEL-08-030390 | RULE ID: SV-230446r627750 | CAT II
5. STIG ID: RHEL-09-654075 | RULE ID: SV-258189r1045352 | CAT II
6. STIG ID: ALMA-09-046660 | RULE ID: SV-269466r1050349 | CAT II
7. STIG ID: UBTU-22-654170 | RULE ID: SV-260636r958446 | CAT II
8. STIG ID: UBTU-24-900350 | RULE ID: SV-270806r1068386 | CAT II

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.	●	●	●

Internal Only - General

6.2.3.34 Ensure the audit configuration is loaded regardless of errors (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Set `-c` flag so that `auditctl` will continue loading rules in spite of an error. The exit code will not be success if any rule fails to load.

Rationale:

The default behaviour of `auditctl` is to stop loading any further rules if it encounters an error in the rules (for example a file watcher referencing a non-existent file). This can lead to auditd running without valid rules being present. It is best to have all valid rules loaded and active rather than a subset.

Audit:

Run the following command and verify output matches:

```
# grep -Ph -- '^\h*-c\b' /etc/audit/rules.d/*.rules | tail -1

-c
```

Remediation:

Edit or create the file `/etc/audit/rules.d/01-initialize.rules` and add the line `-c` at the end of the file:

Example:

```
# printf '%s\n' "" "-c" >> /etc/audit/rules.d/01-initialize.rules
```

Load audit rules

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required.

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000341

Internal Only - General

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 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.		●	●

Internal Only - General

6.2.3.35 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.

Note: This setting will require the system to be rebooted to update the active `auditd` configuration settings.

Rationale:

In immutable mode, unauthorized users cannot execute changes to the audit system to potentially hide malicious activity and then put the audit rules back. Users would most likely notice a system reboot and that could alert administrators of an attempt to make unauthorized audit changes.

Audit:

Run the following command and verify output matches:

```
# grep -Ph -- '^\h*-e\h+2\b' /etc/audit/rules.d/*.rules | tail -1

-e 2
```

Remediation:

Edit or create the file `/etc/audit/rules.d/99-finalize.rules` and add the line `-e 2` at the end of the file:

Example:

```
# printf '\n%s' "-e 2" >> /etc/audit/rules.d/99-finalize.rules
```

Load audit rules

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required.

Internal Only - General

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

References:

1. GRID: NIX-00000340
2. NIST SP 800-53 :: AU-9
3. NIST SP 800-53A :: AU-9.1
4. STIG ID: OL08-00-030121 | RULE ID: SV-248738r958434 | CAT II
5. STIG ID: RHEL-08-030121 | RULE ID: SV-230402r1017208 | CAT II
6. STIG ID: RHEL-09-654275 | RULE ID: SV-258241r1051259 | CAT II
7. STIG ID: ALMA-09-057110 | RULE ID: SV-269195r1050078 | CAT II
8. STIG ID: OL09-00-008005 | RULE ID: SV-271886r1092370 | CAT II
9. STIG ID: UBTU-24-909000 | RULE ID: SV-270832r1068399 | CAT II

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.	●	●	●
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.		●	●

Internal Only - General

6.2.3.36 Ensure the running and on disk configuration is the same (Manual)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

The Audit system has both on disk and running configuration. It is possible for these configuration settings to differ.

Note: Due to the limitations of `augenrules` and `auditctl`, it is not absolutely guaranteed that loading the rule sets via `augenrules --load` will result in all rules being loaded or even that the user will be informed if there was a problem loading the rules.

Rationale:

Configuration differences between what is currently running and what is on disk could cause unexpected problems or may give a false impression of compliance requirements.

Audit:

Merged rule sets

Ensure that all rules in `/etc/audit/rules.d` have been merged into `/etc/audit/audit.rules`:

```
# augenrules --check

/usr/sbin/augenrules: No change
```

Should there be any drift, run `augenrules --load` to merge and load all rules.

Remediation:

If the rules are not aligned across all three (running, on-disk, and merged) areas, run the following command to merge and load all rules:

```
# augenrules --load
```

Check if reboot is required.

Internal Only - General

```
if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then echo "Reboot required
to load rules"; fi
```

References:

1. GRID: NIX-00000343
2. NIST SP 800-53 Rev. 5: AU-3

Additional Information:

Potential reboot required

If the auditing configuration is locked (`-e 2`), then `augenrules` will not warn in any way that rules could not be loaded into the running configuration. A system reboot will be required to load the rules into the running configuration.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.5 Collect Detailed Audit Logs Configure detailed audit logging for enterprise assets containing sensitive data. Include event source, date, username, timestamp, source addresses, destination addresses, and other useful elements that could assist in a forensic investigation.		●	●
v7	6.3 Enable Detailed Logging Enable system logging to include detailed information such as an event source, date, user, timestamp, source addresses, destination addresses, and other useful elements.		●	●

Internal Only - General

6.2.4 Configure auditd File Access

Without the capability to restrict which roles and individuals can select which events are audited, unauthorized personnel may be able to prevent the auditing of critical events.

Internal Only - General

6.2.4.1 Ensure audit log files mode is configured (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Audit log files contain information about the system and system activity.

Rationale:

Access to audit records can reveal system and configuration data to attackers, potentially compromising its confidentiality.

Audit:

Run the following script to verify audit log files are mode `0640` or more restrictive:

Internal Only - General

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

{
   l_perm_mask="0137"
   if [ -e "/etc/audit/auditd.conf" ]; then
      l_audit_log_directory="$(dirname "$(awk -F= '/^\s*log_file\s*/{print
$2}' /etc/audit/auditd.conf | xargs)")"
      if [ -d "$l_audit_log_directory" ]; then
         l_maxperm="$(printf '%o' $(( 0777 & ~$l_perm_mask )) )"
         a_files=()
         while IFS= read -r -d $'\0' l_file; do
            [ -e "$l_file" ] && a_files+=("$l_file")
         done < <(find "$l_audit_log_directory" -maxdepth 1 -type f -perm
/"$l_perm_mask" -print0)
                if (( "${#a_files[@]}" > 0 )); then
            for l_file in "${a_files[@]}"; do
               l_file_mode="$(stat -Lc '%#a' "$l_file")"
               echo -e "\n- Audit Result:\n  ** FAIL **\n  - File:
\"$l_file\" is mode: \"$l_file_mode\"\n     (should be mode: \"$l_maxperm\"
or more restrictive)\n"
            done
         else
            echo -e "\n- Audit Result:\n  ** PASS **\n  - All files in
\"$l_audit_log_directory\" are mode: \"$l_maxperm\" or more restrictive"
         fi
      else
         echo -e "\n- Audit Result:\n  ** FAIL **\n  - Log file directory not
set in \"/etc/audit/auditd.conf\" please set log file directory"
      fi
   else
      echo -e "\n- Audit Result:\n  ** FAIL **\n  - File:
\"/etc/audit/auditd.conf\" not found.\n  - ** Verify auditd is installed **"
   fi
}
```

Remediation:

Run the following command to remove more permissive mode than 0640 from audit log files:

```
# [ -f /etc/audit/auditd.conf ] && find "$(dirname $(awk -F "="
'/^\s*log_file/ {print $2}' /etc/audit/auditd.conf | xargs))" -type f -perm
/0137 -exec chmod u-x,g-wx,o-rwx {} +
```

References:

1. GRID: NIX-00000291
2. NIST SP 800-53 Rev. 5: AU-3

Internal Only - General

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.	●	●	●

Internal Only - General

6.2.4.2 Ensure audit log files owner is configured (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Audit log files contain information about the system and system activity.

Rationale:

Access to audit records can reveal system and configuration data to attackers, potentially compromising its confidentiality.

Audit:

Run the following script to verify audit log files are owned by the `root` user:

Internal Only - General

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

{
   l_output="" l_output2=""
   if [ -e "/etc/audit/auditd.conf" ]; then
      l_audit_log_directory="$(dirname "$(awk -F= '/^\s*log_file\s*/{print
$2}' /etc/audit/auditd.conf | xargs)")"
      if [ -d "$l_audit_log_directory" ]; then
         while IFS= read -r -d $'\0' l_file; do
            l_output2="$l_output2\n  - File: \"$l_file\" is owned by user:
\"$(stat -Lc '%U' "$l_file")\"\n     (should be owned by user: \"root\")\n"
         done < <(find "$l_audit_log_directory" -maxdepth 1 -type f ! -user
root -print0)
      else
         l_output2="$l_output2\n  - Log file directory not set in
\"/etc/audit/auditd.conf\" please set log file directory"
      fi
   else
      l_output2="$l_output2\n  - File: \"/etc/audit/auditd.conf\" not
found.\n  - ** Verify auditd is installed **"
   fi
   if [ -z "$l_output2" ]; then
      l_output="$l_output\n  - All files in \"$l_audit_log_directory\" are
owned by user: \"root\"\n"
      echo -e "\n- Audit Result:\n  ** PASS **\n - * Correctly configured *
:$l_output"
   else
      echo -e "\n- Audit Result:\n  ** FAIL **\n - * Reasons for audit
failure * :$l_output2\n"
   fi
}
```

Remediation:

Run the following command to configure the audit log files to be owned by the `root` user:

```
# [ -f /etc/audit/auditd.conf ] && find "$(dirname $(awk -F "="
'/^\s*log_file/ {print $2}' /etc/audit/auditd.conf | xargs))" -type f ! -user
root -exec chown root {} +
```

References:

1. GRID: NIX-00000292
2. NIST SP 800-53 Rev. 5: AU-3

Internal Only - General

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.	●	●	●

Internal Only - General

6.2.4.3 Ensure audit log files group owner is configured (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Audit log files contain information about the system and system activity.

Rationale:

Access to audit records can reveal system and configuration data to attackers, potentially compromising its confidentiality.

Audit:

Run the following command to verify `log_group` parameter is set to either `adm` or `root` in `/etc/audit/auditd.conf`:

```
#  grep -Piws -- '^\h*log_group\h*=\h*\H+\b' /etc/audit/auditd.conf | grep -
Pvi -- '(adm)'
```

Nothing should be returned

Using the path of the directory containing the audit logs, verify audit log files are owned by the "root" or "adm" group by running the following script:

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

{
  if [ -e /etc/audit/auditd.conf ]; then
      l_fpath="$(dirname "$(awk -F "=" '/^\s*log_file/ {print $2}'
/etc/audit/auditd.conf | xargs)")"
      find -L "$l_fpath" -not -path "$l_fpath"/lost+found -type f \( ! -group
root -a ! -group adm \) -exec ls -l {} +
   fi
}
```

Nothing should be returned

Remediation:

Run the following command to configure the audit log files to be group owned by `adm`:

```
# find $(dirname $(awk -F"=" '/^\s*log_file/ {print $2}'
/etc/audit/auditd.conf | xargs)) -type f \( ! -group adm -a ! -group root \)
-exec chgrp adm {} +
```

Run the following command to set the `log_group` parameter in the audit configuration file to `log_group = adm`:

Internal Only - General

```
# sed -ri 's/^\s*#?\s*log_group\s*=\s*\S+(\s*#.*)?.*$/log_group = adm\1/'
/etc/audit/auditd.conf
```

Run the following command to restart the audit daemon to reload the configuration file:

```
# systemctl restart auditd
```

References:

1. GRID: NIX-00000290
2. NIST SP 800-53 Rev. 5: AU-3

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.	●	●	●

Internal Only - General

6.2.4.4 Ensure the audit log file directory mode is configured (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

The audit log directory contains audit log files.

Rationale:

Audit information includes all information including: audit records, audit settings and audit reports. This information is needed to successfully audit system activity. This information must be protected from unauthorized modification or deletion. If this information were to be compromised, forensic analysis and discovery of the true source of potentially malicious system activity is impossible to achieve.

Audit:

Run the following script to verify the audit log directory is mode 0750 or more restrictive:

Internal Only - General

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

{
   l_perm_mask="0027"
   if [ -e "/etc/audit/auditd.conf" ]; then
      l_audit_log_directory="$(dirname "$(awk -F= '/^\s*log_file\s*/{print
$2}' /etc/audit/auditd.conf | xargs)")"
      if [ -d "$l_audit_log_directory" ]; then
         l_maxperm="$(printf '%o' $(( 0777 & ~$l_perm_mask )) )"
         l_directory_mode="$(stat -Lc '%#a' "$l_audit_log_directory")"
         if [ $(( $l_directory_mode & $l_perm_mask )) -gt 0 ]; then
            echo -e "\n- Audit Result:\n  ** FAIL **\n  - Directory:
\"$l_audit_log_directory\" is mode: \"$l_directory_mode\"\n     (should be
mode: \"$l_maxperm\" or more restrictive)\n"
         else
            echo -e "\n- Audit Result:\n  ** PASS **\n  - Directory:
\"$l_audit_log_directory\" is mode: \"$l_directory_mode\"\n     (should be
mode: \"$l_maxperm\" or more restrictive)\n"
         fi
      else
         echo -e "\n- Audit Result:\n  ** FAIL **\n  - Log file directory not
set in \"/etc/audit/auditd.conf\" please set log file directory"
      fi
   else
      echo -e "\n- Audit Result:\n  ** FAIL **\n  - File:
\"/etc/audit/auditd.conf\" not found\n  - ** Verify auditd is installed **"
   fi
}
```

Remediation:

Run the following command to configure the audit log directory to have a mode of "0750" or less permissive:

```
# chmod g-w,o-rwx "$(dirname "$(awk -F= '/^\s*log_file\s*/{print $2}'
/etc/audit/auditd.conf | xargs)")"
```

Default Value:

750

References:

1. GRID: NIX-00000342
2. NIST SP 800-53 Rev. 5: AU-3

Internal Only - General

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.	●	●	●

Internal Only - General

6.2.4.5 Ensure audit configuration files mode is configured (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Audit configuration files control auditd and what events are audited.

Rationale:

Access to the audit configuration files could allow unauthorized personnel to prevent the auditing of critical events.

Misconfigured audit configuration files may prevent the auditing of critical events or impact the system's performance by overwhelming the audit log. Misconfiguration of the audit configuration files may also make it more difficult to establish and investigate events relating to an incident.

Audit:

Run the following script to verify that the audit configuration files are mode `0640` or more restrictive:

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

{
   l_output="" l_output2="" l_perm_mask="0137"
   l_maxperm="$( printf '%o' $(( 0777 & ~$l_perm_mask )) )"
   while IFS= read -r -d $'\0' l_fname; do
      l_mode=$(stat -Lc '%#a' "$l_fname")
      if [ $(( "$l_mode" & "$l_perm_mask" )) -gt 0 ]; then
         l_output2="$l_output2\n - file: \"$l_fname\" is mode: \"$l_mode\"
(should be mode: \"$l_maxperm\" or more restrictive)"
      fi
   done < <(find /etc/audit/ -type f \( -name "*.conf" -o -name '*.rules' \)
-print0)
   if [ -z "$l_output2" ]; then
      echo -e "\n- Audit Result:\n  ** PASS **\n - All audit configuration
files are mode: \"$l_maxperm\" or more restrictive"
   else
      echo -e "\n- Audit Result:\n  ** FAIL **\n$l_output2"
   fi
}
```

Remediation:

Run the following command to remove more permissive mode than `0640` from the audit configuration files:

Internal Only - General

```
# find /etc/audit/ -type f \( -name '*.conf' -o -name '*.rules' \) -exec
chmod u-x,g-wx,o-rwx {} +
```

References:

1. GRID: NIX-00000288
2. NIST SP 800-53 Rev. 5: AU-3

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.	●	●	●

Internal Only - General

6.2.4.6 Ensure audit configuration files owner is configured (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Audit configuration files control auditd and what events are audited.

Rationale:

Access to the audit configuration files could allow unauthorized personnel to prevent the auditing of critical events.

Misconfigured audit configuration files may prevent the auditing of critical events or impact the system's performance by overwhelming the audit log. Misconfiguration of the audit configuration files may also make it more difficult to establish and investigate events relating to an incident.

Audit:

Run the following command to verify that the audit configuration files are owned by the root user:

```
# find /etc/audit/ -type f \( -name '*.conf' -o -name '*.rules' \) ! -user
root
```

Nothing should be returned

Remediation:

Run the following command to change ownership to `root` user:

```
# find /etc/audit/ -type f \( -name '*.conf' -o -name '*.rules' \) ! -user
root -exec chown root {} +
```

References:

1. GRID: NIX-00000289
2. NIST SP 800-53 Rev. 5: AU-3

Internal Only - General

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.	●	●	●

Internal Only - General

6.2.4.7 Ensure audit configuration files group owner is configured (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Audit configuration files control auditd and what events are audited.

Rationale:

Access to the audit configuration files could allow unauthorized personnel to prevent the auditing of critical events.

Misconfigured audit configuration files may prevent the auditing of critical events or impact the system's performance by overwhelming the audit log. Misconfiguration of the audit configuration files may also make it more difficult to establish and investigate events relating to an incident.

Audit:

Run the following command to verify that the audit configuration files are owned by the group `root`:

```
# find /etc/audit/ -type f \( -name '*.conf' -o -name '*.rules' \) ! -group
root
```

Nothing should be returned

Remediation:

Run the following command to change group to `root`:

```
# find /etc/audit/ -type f \( -name '*.conf' -o -name '*.rules' \) ! -group
root -exec chgrp root {} +
```

References:

1. GRID: NIX-00000287
2. NIST SP 800-53 Rev. 5: AU-3

Internal Only - General

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.	●	●	●

Internal Only - General

6.2.4.8 Ensure audit tools mode is configured (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

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

Rationale:

Protecting audit information includes identifying and protecting the tools used to view and manipulate log data. Protecting audit tools is necessary to prevent unauthorized operation on audit information.

Audit:

Run the following script to verify the audit tools are mode `0755` or more restrictive:

Internal Only - General

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

{
   l_output="" l_output2="" l_perm_mask="0022"
   l_maxperm="$( printf '%o' $(( 0777 & ~$l_perm_mask )) )"
   a_audit_tools=("/sbin/auditctl" "/sbin/aureport" "/sbin/ausearch"
"/sbin/auditd" "/sbin/augenrules")
   for l_audit_tool in "${a_audit_tools[@]}"; do
      l_mode="$(stat -Lc '%#a' "$l_audit_tool")"
      if [ $(( "$l_mode" & "$l_perm_mask" )) -gt 0 ]; then
         l_output2="$l_output2\n - Audit tool \"$l_audit_tool\" is mode:
\"$l_mode\" and should be mode: \"$l_maxperm\" or more restrictive"
      else
         l_output="$l_output\n - Audit tool \"$l_audit_tool\" is correctly
configured to mode: \"$l_mode\""
      fi
   done
   if [ -z "$l_output2" ]; then
      echo -e "\n- Audit Result:\n  ** PASS **\n - * Correctly configured *
:$l_output"
   else
      echo -e "\n- Audit Result:\n  ** FAIL **\n - * Reasons for audit
failure * :$l_output2\n"
      [ -n "$l_output" ] && echo -e "\n - * Correctly configured *
:\n$l_output\n"
   fi
   unset a_audit_tools
}
```

Remediation:

Run the following command to remove more permissive mode from the audit tools:

```
# chmod go-w /sbin/auditctl /sbin/aureport /sbin/ausearch /sbin/auditd
/sbin/augenrules
```

References:

1. GRID: NIX-00000297
2. NIST SP 800-53 Rev. 5: AU-3

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.	●	●	●

Internal Only - General

Controls Version	Control	IG 1	IG 2	IG 3
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.	●	●	●

Internal Only - General

6.2.4.9 Ensure audit tools owner is configured (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

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

Rationale:

Protecting audit information includes identifying and protecting the tools used to view and manipulate log data. Protecting audit tools is necessary to prevent unauthorized operation on audit information.

Audit:

Run the following command to verify the audit tools are owned by the `root` user:

```
# stat -Lc "%n %U" /sbin/auditctl /sbin/aureport /sbin/ausearch /sbin/auditd
/sbin/augenrules | awk '$2 != "root" {print}'
```

Nothing should be returned

Remediation:

Run the following command to change the owner of the audit tools to the `root` user:

```
# chown root /sbin/auditctl /sbin/aureport /sbin/ausearch /sbin/auditd
/sbin/augenrules
```

References:

1. GRID: NIX-00000298
2. NIST SP 800-53 Rev. 5: AU-3

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.	●	●	●

Internal Only - General

Controls Version	Control	IG 1	IG 2	IG 3
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.	●	●	●

Internal Only - General

6.2.4.10 Ensure audit tools group owner is configured (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

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

Rationale:

Protecting audit information includes identifying and protecting the tools used to view and manipulate log data. Protecting audit tools is necessary to prevent unauthorized operation on audit information.

Audit:

Run the following command to verify the audit tools are owned by the group `root`

```
# stat -Lc "%n %G" /sbin/auditctl /sbin/aureport /sbin/ausearch /sbin/auditd
/sbin/augenrules | awk '$2 != "root" {print}'
```

Nothing should be returned

Remediation:

Run the following command to change group ownership to the group `root`:

```
# chgrp root /sbin/auditctl /sbin/aureport /sbin/ausearch /sbin/auditd
/sbin/augenrules
```

References:

1. GRID: NIX-00000296
2. NIST SP 800-53 Rev. 5: AU-3

Internal Only - General

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.	●	●	●

Internal Only - General

6.3 Configure 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.

Internal Only - General

6.3.1 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 to verify `aide` is installed:

```
# dpkg-query -s aide &>/dev/null && echo "aide is installed"

aide is installed
```

Run the following command to verify `aide-common` is installed:

```
# dpkg-query -s aide-common &>/dev/null && echo "aide-common is installed"

aide-common is installed
```

Remediation:

Install AIDE using the appropriate package manager or manual installation:

```
# apt install aide aide-common
```

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

Run the following commands to initialize AIDE:

```
# aideinit
# mv /var/lib/aide/aide.db.new /var/lib/aide/aide.db
```

References:

1. GRID: NIX-00000347
2. NIST SP 800-53 Rev. 5: AU-2
3. STIG ID: UBTU-20-010450 | Rule ID: SV-238371r958944 | CAT II
4. STIG ID: UBTU-22-651010 | Rule ID: SV-260582r958944 | CAT II

Internal Only - General

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

Internal Only - General

6.3.2 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 command:

```
# systemctl list-unit-files | awk
'$1~/^dailyaidecheck\.(timer|service)$/{print $1 "\t" $2}'
```

Example output:

```
dailyaidecheck.service  static
dailyaidecheck.timer    enabled
```

Verify `dailyaidecheck.timer` is `enabled` and `dailyaidecheck.service` is either `static` or `enabled`.

Run the following command to verify `dailyaidecheck.timer` is `active`:

```
# systemctl is-active dailyaidecheck.timer

active
```

Remediation:

Run the following command to unmask `dailyaidecheck.timer` and `dailyaidecheck.service`:

```
# systemctl unmask dailyaidecheck.timer dailyaidecheck.service
```

Run the following command to enable and start `dailyaidecheck.timer`:

```
# systemctl --now enable dailyaidecheck.timer
```

References:

1. GRID: NIX-00000349
2. https://github.com/konstruktoid/hardening/blob/master/config/aidecheck.service

Internal Only - General

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

Additional Information:

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

systemd timers, timer file `aidecheck.timer` and service file `aidecheck.service`, have been included as an optional alternative to using `cron`

Ubuntu advises using `/usr/bin/aide.wrapper` rather than calling `/usr/bin/aide` directly in order to protect the database and prevent conflicts

CIS Controls:

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

Internal Only - General

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

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

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

`aide.conf` is case-sensitive. Leading and trailing white spaces are ignored. Each config lines must end with new line.

`AIDE` uses the backslash character `\` as escape character for ' ' (space), '@' and '' (backslash) (e.g. '\ ' or '@'). To literally match a '' in a file path with a regular expression you have to escape the backslash twice (i.e. '\\').

There are three types of lines in `aide.conf`:

- The configuration options which are used to set configuration parameters and define groups.
- (restricted) rules that are used to indicate which files are added to the database.
- Macro lines define or undefine variables within the config file.

Note: Lines beginning with `#` are ignored as comments.

`@@include` <*FILE*> - Include <*FILE*>.

- The content of the file is used as if it were inserted in this part of the config file.
- The maximum depth of nested includes is 16.

`@@include <*DIRECTORY*> <*REGEX*> - [RULE_PREFIX] (added in AIDE v0.17)

- Include all (regular) files found in <DIRECTORY> matching regular expression <REGEX> (sub-directories are ignored).
- The file are included in lexical sort order.
- If `RULE_PREFIX` is set, all rules included by the statement are prefixed with given <RULE_PREFIX> (added in AIDE v0.18). Prefixes from nested include statements are concatenated.
- The content of the files is used as if it were inserted in this part of the config file.

`@x_include`:

Internal Only - General

- is identical to `@@include`, except that if a config file is executable is is run and the output is used as config.
- If the executable file exits with status greater than zero or writes to stderr aide stops with an error.
- For security reasons <*DIRECTORY*> and each executable config file must be owned by the current user or root. They must not be group- or world-writable.
- `@@x_include _<FILE>_` (added in AIDE v0.17):
- `@@x_include <*DIRECTORY*> <*REGEX*> [RULE_PREFIX] (added in AIDE v0.17)

`@@x_include_setenv` <*VAR*> <*VALUE*> (added in AIDE v0.17)

- Adds the variable <*VAR*> with the value <*VALUE*> to the environment used for config file execution.
- Environment variable names are limited to alphanumeric characters (A-Za-z0-9) and the underscore '_' and must not begin with a digit.

Rationale:

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

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

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

Audit:

Verify that Advanced Intrusion Detection Environment (AIDE) is properly configured .

Run the following script to verify:

- AIDE is configured to use cryptographic mechanisms to protect the integrity of audit tools:
- The following audit tool files include the options "p, i, n, u, g, s, b, acl, xattrs and sha512"
    - auditctl
    - auditd
    - ausearch
    - aureport
    - augenrules

Internal Only - General

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

{
   a_output=() a_output2=() l_tool_dir="$(readlink -f /sbin)"
   a_items=("p" "i" "n" "u" "g" "s" "b" "acl" "xattrs" "sha512")
   l_aide_cmd="$(whereis aide | awk '{print $2}')"
   a_audit_files=("auditctl" "auditd" "ausearch" "aureport" "augenrules")
   if [ -f "$l_aide_cmd" ] && command -v "$l_aide_cmd" &>/dev/null; then
      a_aide_conf_files=("$(find -L /etc -type f -name 'aide.conf')")
      f_file_par_chk()
      {
         a_out2=()
         for l_item in "${a_items[@]}"; do
            ! grep -Psiq -- '(\h+|\+)'"$l_item"'(\h+|\+)' <<< "$l_out" && \
            a_out2+=("  - Missing the \"$l_item\" option")
         done
         if [ "${#a_out2[@]}" -gt "0" ]; then
            a_output2+=(" - Audit tool file: \"$l_file\"" "${a_out2[@]}")
         else
            a_output+=(" - Audit tool file: \"$l_file\" includes:" "
\"${a_items[*]}\"")
         fi
      }
      for l_file in "${a_audit_files[@]}"; do
         if [ -f "$l_tool_dir/$l_file" ]; then
            l_out="$("$l_aide_cmd" --config "${a_aide_conf_files[@]}" -p
f:"$l_tool_dir/$l_file")"
            f_file_par_chk
         else
            a_output+=("  - Audit tool file \"$l_file\" doesn't exist")
         fi
      done
   else
      a_output2+=("  - The command \"aide\" was not found"  "    Please
install AIDE")
   fi
   if [ "${#a_output2[@]}" -le 0 ]; then
      printf '%s\n' "" "- Audit Result:" "  ** PASS **" "${a_output[@]}" ""
   else
      printf '%s\n' "" "- Audit Result:" "  ** FAIL **" " - Reason(s) for
audit failure:" "${a_output2[@]}"
      [ "${#a_output[@]}" -gt 0 ] && printf '%s\n' "" "- Correctly set:"
"${a_output[@]}" ""
   fi
}
```

Note: The script is written to read the "winning" configuration setting, to include any configuration settings in files included as part of the `@@x_include` setting.

Remediation:

Run the following command to determine the absolute path to the non-symlinked version on the audit tools:

```
# readlink -f /sbin
```

The output will be either `/usr/sbin` **- OR -** `/sbin`. Ensure the correct path is used.

Internal Only - General

Edit /etc/aide/aide.conf and add or update the following selection lines replacing <PATH> with the correct path returned in the command above:

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

Example

```
# printf '%s\n' "" "# Audit Tools" "$(readlink -f /sbin/auditctl)
p+i+n+u+g+s+b+acl+xattrs+sha512" "$(readlink -f /sbin/auditd)
p+i+n+u+g+s+b+acl+xattrs+sha512" "$(readlink -f /sbin/ausearch)
p+i+n+u+g+s+b+acl+xattrs+sha512" "$(readlink -f /sbin/aureport)
p+i+n+u+g+s+b+acl+xattrs+sha512" "$(readlink -f /sbin/augenrules)
p+i+n+u+g+s+b+acl+xattrs+sha512" >> /etc/aide/aide.conf
```

Note: - IF - /etc/aide/aide.conf includes a @@x_include statement:

- *<DIRECTORY>* and each executable config file must be owned by the current user or root
- They must not be group or world-writable

Example:

```
@@x_include /etc/aide.conf.d ^[a-zA-Z0-9_-]+$
```

References:

1. GRID: NIX-00000348
2. AIDE.CONF(5)

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

7 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.

Internal Only - General

7.1 Configure system file and directory access

This section provides guidance on securing aspects of system files and directories.

Internal Only - General

7.1.1 Ensure access to /etc/passwd is 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 to verify `/etc/passwd` is mode 644 or more restrictive, `Uid` is `0/root` and `Gid` is `0/root`:

```
# stat -Lc 'Access: (%#a/%A)  Uid: ( %u/ %U) Gid: ( %g/ %G)'  /etc/passwd

Access: (0644/-rw-r--r--)  Uid: ( 0/ root) Gid: ( 0/ root)
```

Remediation:

Run the following commands to remove excess permissions, set owner, and set group on `/etc/passwd`:

```
# chmod u-x,go-wx /etc/passwd
# chown root:root /etc/passwd
```

Default Value:

Access: (0644/-rw-r--r--) Uid: (0/ root) Gid: (0/ root)

References:

1. GRID: NIX-00000354
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

Internal Only - General

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.	●	●	●

Internal Only - General

7.1.2 Ensure access to /etc/passwd- is 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 to verify `/etc/passwd-` is mode 644 or more restrictive, `Uid` is `0/root` and `Gid` is `0/root`:

```
# stat -Lc 'Access: (%#a/%A)  Uid: ( %u/ %U) Gid: { %g/ %G)' /etc/passwd-

Access: (0644/-rw-r--r--)  Uid: ( 0/ root) Gid: { 0/ root)
```

Remediation:

Run the following commands to remove excess permissions, set owner, and set group on `/etc/passwd-`:

```
# chmod u-x,go-wx /etc/passwd-
# chown root:root /etc/passwd-
```

Default Value:

Access: (0644/-rw-r--r--) Uid: (0/ root) Gid: { 0/ root)

References:

1. GRID: NIX-00000355
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

Internal Only - General

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.	●	●	●

Internal Only - General

7.1.3 Ensure access to /etc/group is 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 to verify `/etc/group` is mode 644 or more restrictive, `Uid` is `0/root` and `Gid` is `0/root`:

```
# stat -Lc 'Access: (%#a/%A)  Uid: ( %u/ %U) Gid: ( %g/ %G)' /etc/group

Access: (0644/-rw-r--r--)  Uid: ( 0/ root) Gid: ( 0/ root)
```

Remediation:

Run the following commands to remove excess permissions, set owner, and set group on `/etc/group`:

```
# chmod u-x,go-wx /etc/group
# chown root:root /etc/group
```

Default Value:

Access: (0644/-rw-r--r--) Uid: (0/ root) Gid: (0/ root)

References:

1. GRID: NIX-00000350
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

Internal Only - General

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.	●	●	●

Internal Only - General

7.1.4 Ensure access to /etc/group- is 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 to verify `/etc/group-` is mode 644 or more restrictive, `Uid` is `0/root` and `Gid` is `0/root`:

```
# stat -Lc 'Access: (%#a/%A)  Uid: ( %u/ %U) Gid: ( %g/ %G)'  /etc/group-

Access: (0644/-rw-r--r--)  Uid: ( 0/ root) Gid: ( 0/ root)
```

Remediation:

Run the following commands to remove excess permissions, set owner, and set group on `/etc/group-`:

```
# chmod u-x,go-wx /etc/group-
# chown root:root /etc/group-
```

Default Value:

Access: (0644/-rw-r--r--) Uid: (0/ root) Gid: (0/ root)

References:

1. GRID: NIX-00000351
2. NIST SP 800-53 Rev. 5: AC-3. MP-2

Internal Only - General

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.	●	●	●

Internal Only - General

7.1.5 Ensure access to /etc/shadow is 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 to verify `/etc/shadow` is mode 640 or more restrictive, `Uid` is `0/root` and `Gid` is `0/root` or ({GID}/ shadow):

```
# stat -Lc 'Access: (%#a/%A)  Uid: ( %u/ %U) Gid: ( %g/ %G)'  /etc/shadow
```

Example:

```
Access: (0640/-rw-r-----)  Uid: ( 0/ root) Gid: ( 42/ shadow)
```

Remediation:

Run **one** of the following commands to set ownership of `/etc/shadow` to `root` and group to either `root` or `shadow`:

```
# chown root:shadow /etc/shadow
  -OR-
# chown root:root /etc/shadow
```

Run the following command to remove excess permissions form `/etc/shadow`:

```
# chmod u-x,g-wx,o-rwx /etc/shadow
```

Default Value:

Access: (0640/-rw-r-----) Uid: (0/ root) Gid: (42/ shadow)

References:

1. GRID: NIX-00000357
2. NIST SP 800-53 Rev. 5: AC-3. MP-2

Internal Only - General

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.	●	●	●

Internal Only - General

7.1.6 Ensure access to /etc/shadow- is 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 to verify `/etc/shadow-` is mode 640 or more restrictive, `Uid` is `0/root` and `Gid` is `0/root` or `{GID}/shadow`:

```
# stat -Lc 'Access: (%#a/%A)  Uid: ( %u/ %U) Gid: ( %g/ %G)'  /etc/shadow-
```

Example:

```
Access: (0640/-rw-r-----)  Uid: ( 0/ root) Gid: ( 42/ shadow)
```

Remediation:

Run **one** of the following commands to set ownership of `/etc/shadow-` to `root` and group to either `root` or `shadow`:

```
# chown root:shadow /etc/shadow-
  -OR-
# chown root:root /etc/shadow-
```

Run the following command to remove excess permissions form `/etc/shadow-`:

```
# chmod u-x,g-wx,o-rwx /etc/shadow-
```

Default Value:

Access: (0640/-rw-r-----) Uid: (0/ root) Gid: (42/ shadow)

References:

1. GRID: NIX-00000358
2. NIST SP 800-53 Rev. 5: AC-3. MP-2

Internal Only - General

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.	●	●	●

Internal Only - General

7.1.7 Ensure access to /etc/gshadow is 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 to verify `/etc/gshadow` is mode 640 or more restrictive, `Uid` is `0/root` and `Gid` is `0/root` or `{GID}/shadow`:

```
# stat -Lc 'Access: (%#a/%A)  Uid: ( %u/ %U) Gid: ( %g/ %G)'  /etc/gshadow
```

Example:

```
Access: (0640/-rw-r-----)  Uid: ( 0/ root) Gid: ( 42/ shadow)
```

Note: `/etc/gshadow` may not exist by default. Security best practice recommends creating `/etc/gshadow` to ensure hashed group passwords are not stored in `/etc/group`

Remediation:

- IF - `/etc/gshadow` does not exist run the following command to convert group entries and create `/etc/gshadow` if needed:

```
# grpconv
```

Run **one** of the following commands to set ownership of `/etc/gshadow` to `root` and group to either `root` or `shadow`:

```
# chown root:shadow /etc/gshadow
  -OR-
# chown root:root /etc/gshadow
```

Run the following command to remove excess permissions form `/etc/gshadow`:

Internal Only - General

```
# chmod u-x,g-wx,o-rwx /etc/gshadow
```

Default Value:

Access: (0640/-rw-r-----) Uid: (0/ root) Gid: (42/ shadow)

References:

1. GRID: NIX-00000352
2. NIST SP 800-53 Rev. 5: AC-3. MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 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.	●	●	●

Internal Only - General

7.1.8 Ensure access to /etc/gshadow- is 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 to verify `/etc/gshadow-` is mode 640 or more restrictive, `Uid` is `0/root` and `Gid` is `0/root` or `{GID}/shadow`:

```
# stat -Lc 'Access: (%#a/%A)  Uid: ( %u/ %U) Gid: ( %g/ %G)'  /etc/gshadow-
```

Example:

```
Access: (0640/-rw-r-----)  Uid: ( 0/ root) Gid: ( 42/ shadow)
```

Note: If `/etc/gshadow-` does not exist this is considered a passing state

Remediation:

Run **one** of the following commands to set ownership of `/etc/gshadow-` to `root` and group to either `root` or `shadow`:

```
# chown root:shadow /etc/gshadow-
  -OR-
# chown root:root /etc/gshadow-
```

Run the following command to remove excess permissions form `/etc/gshadow-`:

```
# chmod u-x,g-wx,o-rwx /etc/gshadow-
```

Default Value:

Access: (0640/-rw-r-----) Uid: (0/ root) Gid: (42/ shadow)

References:

1. GRID: NIX-00000353
2. NIST SP 800-53 Rev. 5: AC-3. MP-2

Internal Only - General

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.	●	●	●

Internal Only - General

7.1.9 Ensure access to /etc/shells is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`/etc/shells` is a text file which contains the full pathnames of valid login shells. This file is consulted by `chsh` and available to be queried by other programs.

Rationale:

It is critical to ensure that the `/etc/shells` file is protected from unauthorized access. Although it is protected by default, the file permissions could be changed either inadvertently or through malicious actions.

Audit:

Run the following command to verify `/etc/shells` is mode 644 or more restrictive, `Uid` is `0/root` and `Gid` is `0/root`:

```
# stat -Lc 'Access: (%#a/%A)  Uid: ( %u/ %U) Gid: ( %g/ %G)' /etc/shells

Access: (0644/-rw-r--r--)  Uid: ( 0/ root) Gid: ( 0/ root)
```

Remediation:

Run the following commands to remove excess permissions, set owner, and set group on `/etc/shells`:

```
# chmod u-x,go-wx /etc/shells
# chown root:root /etc/shells
```

Default Value:

Access: (0644/-rw-r--r--) Uid: (0/ root) Gid: (0/ root)

References:

1. GRID: NIX-00000359
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

Internal Only - General

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.	●	●	●

Internal Only - General

7.1.10 Ensure access to /etc/security/opasswd is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`/etc/security/opasswd` and it's backup `/etc/security/opasswd.old` hold user's previous passwords if `pam_unix` or `pam_pwhistory` is in use on the system

Rationale:

It is critical to ensure that `/etc/security/opasswd` 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 commands to verify `/etc/security/opasswd` and `/etc/security/opasswd.old` are mode 600 or more restrictive, `Uid` is `0/root` and `Gid` is `0/root` if they exist:

```
#  [ -e "/etc/security/opasswd" ] && stat -Lc '%n Access: (%#a/%A)  Uid: (
%u/ %U) Gid: ( %g/ %G)' /etc/security/opasswd

/etc/security/opasswd Access: (0600/-rw-------)  Uid: ( 0/ root) Gid: ( 0/
root)
 -OR-
Nothing is returned
#  [ -e "/etc/security/opasswd.old" ] && stat -Lc '%n Access: (%#a/%A)  Uid:
( %u/ %U) Gid: ( %g/ %G)' /etc/security/opasswd.old

/etc/security/opasswd.old Access: (0600/-rw-------)  Uid: ( 0/ root) Gid: (
0/ root)
 -OR-
Nothing is returned
```

Remediation:

Run the following commands to remove excess permissions, set owner, and set group on `/etc/security/opasswd` and `/etc/security/opasswd.old` is they exist:

Internal Only - General

```
# [ -e "/etc/security/opasswd" ] && chmod u-x,go-rwx /etc/security/opasswd
# [ -e "/etc/security/opasswd" ] && chown root:root /etc/security/opasswd
# [ -e "/etc/security/opasswd.old" ] && chmod u-x,go-rwx
/etc/security/opasswd.old
# [ -e "/etc/security/opasswd.old" ] && chown root:root
/etc/security/opasswd.old
```

Default Value:

/etc/security/opasswd Access: (0600/-rw-------) Uid: (0/ root) Gid: (0/ root)

References:

1. GRID: NIX-00000356
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 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.	●	●	●

Internal Only - General

7.1.11 Ensure world writable files and directories are secured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

World writable files are the least secure. Data in world-writable files can be modified and compromised by any user on the system. World writable files may also indicate an incorrectly written script or program that could potentially be the cause of a larger compromise to the system's integrity. See the `chmod(2)` man page for more information.

Setting the sticky bit on world writable directories prevents users from deleting or renaming files in that directory that are not owned by them.

Rationale:

Data in world-writable files can be modified and compromised by any user on the system. World writable files may also indicate an incorrectly written script or program that could potentially be the cause of a larger compromise to the system's integrity.

This feature prevents the ability to delete or rename files in world writable directories (such as `/tmp`) that are owned by another user.

Audit:

Run the following script to verify:

- No world writable files exist
- No world writable directories without the sticky bit exist

Internal Only - General

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

{
   a_output=() a_output2=() a_file=() a_dir=() l_smask='01000'
   l_exclude="nfs|proc|cifs|smb|vfat|iso9660|efivarfs|selinuxfs|ncpfs"
   l_exclude2="\/run|\/tmp|\/var\/tmp"
   a_path=(-path "*/containers/storage/*" -o -path "*/containerd/*" -o -path
"*/kubelet/*" -o -path "/sys/*" -o -path "/snap/*" -o -path "/boot/efi/*")
   while IFS= read -r l_mount; do
      while IFS= read -r -d $'\0' l_file; do
         if [ -e "$l_file" ]; then
            [ -f "$l_file" ] && a_file+=("$l_file")
            if [ -d "$l_file" ]; then
               l_mode="$(stat -Lc '%#a' "$l_file")"
               [ ! $(( $l_mode & $l_smask )) -gt 0 ] && a_dir+=("$l_file")
            fi
         fi
      done < <(find "$l_mount" -mount -xdev \( "${a_path[@]}" \) \
      -prune -o \( -type f -o -type d \) -perm -0002 -print0 2> /dev/null)
   done < <(findmnt -Dkerno fstype,target | \
   awk '($1 !~ /^\s*('"$l_exclude"')/ && $2 !~
/^('"$l_exclude2"')(\/|$)/){print $2}')
   if [ "${#a_file[@]}" -le 0 ]; then
      printf '%s\n' "" "  - No world writable files exist on the local
filesystem."
   else
      printf '%s\n' "" " - There are \"${#a_file[@]}\" World writable files
on the system." \
      "   - The following is a list of World writable files:" \
      "${a_file[@]}" "   - end of list"
   fi
   if [ "${#a_dir[@]}" -le 0 ]; then
      printf '%s\n' "" \
      "  - Sticky bit is set on world writable directories on the local
filesystem."
   else
      printf '%s\n' "" " - There are \"${#a_dir[@]}\" World writable
directories" \
      "    without the sticky bit on the system." \
      "   - The following is a list of World writable directories without the
sticky bit:" \
      "${a_dir[@]}" "   - end of list"
   fi
}
```

Note: On systems with a large number of files and/or directories, this audit may be a long running process

Remediation:

- World Writable Files:
  - It is recommended that write access is removed from `other` with the command (`chmod o-w <filename>`), but always consult relevant vendor documentation to avoid breaking any application dependencies on a given file.

Internal Only - General

- World Writable Directories:
  - Set the sticky bit on all world writable directories with the command (`chmod a+t <directory_name>`)

Run the following script to:

- Remove other write permission from any world writable files
- Add the sticky bit to all world writable directories

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

{
   l_smask='01000' a_file=() a_dir=()
   a_path=(-path "*/containers/storage/*" -o -path "*/containerd/*" -o -path
"*/kubelet/*" -o -path "/sys/*" -o -path "/snap/*" -o -path "/boot/efi/*")
   l_exclude="nfs|proc|cifs|smb|vfat|iso9660|efivarfs|selinuxfs|ncpfs"
   l_exclude2="\/run|\/tmp|\/var\/tmp"
   while IFS= read -r l_mount; do
      while IFS= read -r -d $'\0' l_file; do
         l_mode="$(stat -Lc '%#a' "$l_file" 2>/dev/null)"
         if [ -f "$l_file" ]; then # Remove excess permissions from WW files
            printf '%s\n' "" " - File: \"$l_file\" is mode: \"$l_mode\"" \
            "  - removing write permission on \"$l_file\" from \"other\""
            chmod o-w "$l_file"
         fi
         if [ -d "$l_file" ]; then # Add sticky bit
            if [ ! $(( $l_mode & $l_smask )) -gt 0 ]; then
               printf '%s\n' "" " - Directory: \"$l_file\" is mode:
\"$l_mode\"" \
               "    and doesn't have the sticky bit set" \
               "  - Adding the sticky bit"
               chmod a+t "$l_file"
            fi
         fi
      done < <(find "$l_mount" -mount -xdev \( "${a_path[@]}" \) \
      -prune -o \( -type f -o -type d \) -perm -0002 -print0 2> /dev/null)
   done < <(findmnt -Dkerno fstype,target | \
   awk '($1 !~ /^\s*('"$l_exclude"')/ && $2 !~
/^('"$l_exclude2"')(\/|$)/){print $2}')
}
```

References:

1. GRID: NIX-00000362
2. NIST SP 800-53 Rev. 5: AC-3. MP-2
3. STIG ID: UBTU-20-010411 | Rule ID: SV-238332r958524 | CAT ||
4. STIG ID: UBTU-22-232145 | Rule ID: SV-260513r958524 | CAT ||

Internal Only - General

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.	●	●	●

Internal Only - General

7.1.12 Ensure no files or directories without an owner and a group exist (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Administrators may delete users or groups from the system and neglect to remove all files and/or directories owned by those users or groups.

Rationale:

A new user or group who is assigned a deleted user's user ID or group ID may then end up "owning" a deleted user or group's files, and thus have more access on the system than was intended.

Audit:

Run the following script to verify no unowned or ungrouped files or directories exist:

Internal Only - General

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

{
   a_output=() a_output2=() a_nouser=() a_nogroup=() # Initialize arrays
   a_path=(-path "*/containers/storage/*" -o -path "*/containerd/*" -o -path
"*/kubelet/*" -o -path "/sys/*" -o -path "/snap/*" -o -path "/boot/efi/*")
   while IFS= read -r l_mount; do
      while IFS= read -r -d $'\0' l_file; do
         if [ -e "$l_file" ]; then
            while IFS=: read -r l_user l_group; do
               [ "$l_user" = "UNKNOWN" ] && a_nouser+=("$l_file")
               [ "$l_group" = "UNKNOWN" ] && a_nogroup+=("$l_file")
            done < <(stat -Lc '%U:%G' "$l_file")
         fi
      done < <(find "$l_mount" -mount -xdev \( "${a_path[@]}" \) -prune -o \(
-type f -o -type d \) \( -nouser -o -nogroup \) -print0 2> /dev/null)
   done < <(findmnt -Dkerno fstype,target | awk '($1 !~
/^\s*(nfs|proc|cifs|smb|vfat|iso9660|efivarfs|selinuxfs|ncpfs)/ && $2 !~
/^(\/run|\/tmp|\/var\/tmp)(\/|$)/){print $2}')
   if [ "${#a_nouser[@]}" -le 0 ]; then
      a_output+=("  - No files or directories without an owner exist on the
local filesystem.")
   else
      a_output2+=("  - There are \"${#a_nouser[@]}\" unowned files or
directories on the system." \
      "   - The following is a list of unowned files and/or directories:" \
      "${a_nouser[@]}" "   - end of list")
   fi
   if [ "${#a_nogroup[@]}" -le 0 ]; then
      a_output+=("  - No files or directories without a group owner exist on
the local filesystem.")
   else
      a_output2+=("  - There are \"${#a_nogroup[@]}\" ungrouped files or
directories on the system." \
      "   - The following is a list of ungrouped files and/or directories:" \
      "${a_nogroup[@]}" "   - end of list")
   fi
   if [ "${#a_output2[@]}" -le 0 ]; then
      printf '%s\n' "" "- Audit Result:" "  ** PASS **" "${a_output[@]}"
   else
      printf '%s\n' "" "- Audit Result:" "  ** FAIL **" "  * Reasons for
audit failure *" "${a_output2[@]}" ""
      [ "${#a_output[@]}" -gt 0 ] && printf '%s\n' "- Correctly set:"
"${a_output[@]}"
   fi
}
```

Note: On systems with a large number of files and/or directories, this audit may be a long running process

Remediation:

Remove or set ownership and group ownership of these files and/or directories to an active user on the system as appropriate.

Internal Only - General

References:

1. GRID: NIX-00000360
2. NIST SP 800-53 Rev. 5: AC-3. MP-2

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.3 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.	●	●	●

Internal Only - General

7.1.13 Ensure SUID and SGID files are reviewed (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 or SGID program is to enable users to perform functions (such as changing their password) that require root privileges.

Rationale:

There are valid reasons for SUID and SGID programs, but it is important to identify and review such programs to ensure they are legitimate. Review the files returned by the action in the audit section and check to see if system binaries have a different checksum than what from the package. This is an indication that the binary may have been replaced.

Audit:

Run the following script to generate a list of SUID and SGID files:

Internal Only - General

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

{
   l_output="" l_output2=""
   a_suid=(); a_sgid=() # initialize arrays
   while IFS= read -r l_mount; do
      while  IFS= read -r -d $'\0' l_file; do
         if [ -e "$l_file" ]; then
            l_mode="$(stat -Lc '%#a' "$l_file")"
            [ $(( $l_mode & 04000 )) -gt 0 ] && a_suid+=("$l_file")
            [ $(( $l_mode & 02000 )) -gt 0 ] && a_sgid+=("$l_file")
         fi
      done < <(find "$l_mount" -xdev -type f \( -perm -2000 -o -perm -4000 \)
-print0 2>/dev/null)
   done < <(findmnt -Dkerno fstype,target,options | awk '($1 !~
/^\s*(nfs|proc|smb|vfat|iso9660|efivarfs|selinuxfs)/ && $2 !~
/^\/run\/user\// && $3 !~/noexec/ && $3 !~/nosuid/) {print $2}')
   if ! (( ${#a_suid[@]} > 0 )); then
      l_output="$l_output\n - No executable SUID files exist on the system"
   else
      l_output2="$l_output2\n - List of \"$(printf '%s' "${#a_suid[@]}")\"
SUID executable files:\n$(printf '%s\n' "${a_suid[@]}")\n - end of list -\n"
   fi
   if ! (( ${#a_sgid[@]} > 0 )); then
      l_output="$l_output\n - No SGID files exist on the system"
   else
      l_output2="$l_output2\n - List of \"$(printf '%s' "${#a_sgid[@]}")\"
SGID executable files:\n$(printf '%s\n' "${a_sgid[@]}")\n - end of list -\n"
   fi
   [ -n "$l_output2" ] && l_output2="$l_output2\n- Review the preceding
list(s) of SUID and/or SGID files to\n- ensure that no rogue programs have
been introduced onto the system.\n"
   unset a_arr; unset a_suid; unset a_sgid # Remove arrays
   # If l_output2 is empty, Nothing to report
   if [ -z "$l_output2" ]; then
      echo -e "\n- Audit Result:\n$l_output\n"
   else
      echo -e "\n- Audit Result:\n$l_output2\n"
      [ -n "$l_output" ] && echo -e "$l_output\n"
   fi
}
```

Note: on systems with a large number of files, this may be a long running process

Remediation:

Ensure that no rogue SUID or SGID programs have been introduced into the system. Review the files returned by the action in the Audit section and confirm the integrity of these binaries.

References:

1. GRID: NIX-00000361
2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5, AC-3, MP-2

Internal Only - General

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.	●	●	●

Internal Only - General

7.2 Local User and Group Settings

This section provides guidance on securing aspects of the local users and groups.

Note: The recommendations in this section check local users and groups. Any users or groups from other sources such as LDAP will not be audited. In a domain environment similar checks should be performed against domain users and groups.

Internal Only - General

7.2.1 Ensure accounts in /etc/passwd use shadowed passwords (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Local accounts can use shadowed passwords. With shadowed passwords, the passwords are saved in the shadow password file, `/etc/shadow`, encrypted by a salted one-way hash. Accounts with a shadowed password have an `x` in the second field in `/etc/passwd`.

Rationale:

The `/etc/passwd` file also contains information like user ID's and group ID's that are used by many system programs. Therefore, the `/etc/passwd` file must remain world readable. In spite of encoding the password with a randomly-generated one-way hash function, an attacker could still break the system if they got access to the `/etc/passwd` file. This can be mitigated by using shadowed passwords, thus moving the passwords in the `/etc/passwd` file to `/etc/shadow`. The `/etc/shadow` file is set so only root will be able to read and write. This helps mitigate the risk of an attacker gaining access to the encoded passwords with which to perform a dictionary attack.

Note:

- All accounts must have passwords or be locked to prevent the account from being used by an unauthorized user.
- A user account with an empty second field in `/etc/passwd` allows the account to be logged into by providing only the username.

Audit:

Run the following command and verify that no output is returned:

```
# awk -F: '($2 != "x" ) { print "User: \"" $1 "\" is not set to shadowed passwords"}' /etc/passwd
```

Remediation:

Run the following command to set accounts to use shadowed passwords and migrate passwords in `/etc/passwd` to `/etc/shadow`:

```
# pwconv
```

Investigate to determine if the account is logged in and what it is being used for, to determine if it needs to be forced off.

Internal Only - General

References:

1. GRID: NIX-00000364
2. NIST SP 800-53 Rev. 5: IA-5
3. PWCONV(8)

Additional Information:

The `pwconv` command creates shadow from `passwd` and an optionally existing `shadow`.

- The `pwunconv` command creates `passwd` from `passwd` and `shadow` and then removes `shadow`.
- The `grpconv` command creates `gshadow` from `group` and an optionally existing `gshadow`.
- The `grpunconv` command creates `group` from `group` and `gshadow` and then removes `gshadow`.

These four programs all operate on the normal and shadow password and group files: `/etc/passwd`, `/etc/group`, `/etc/shadow`, and `/etc/gshadow`.

Each program acquires the necessary locks before conversion. `pwconv` and `grpconv` are similar. First, entries in the shadowed file which don't exist in the main file are removed. Then, shadowed entries which don't have `x' as the password in the main file are updated. Any missing shadowed entries are added. Finally, passwords in the main file are replaced with` x'. These programs can be used for initial conversion as well to update the shadowed file if the main file is edited by hand.

`pwconv` will use the values of `PASS_MIN_DAYS`, `PASS_MAX_DAYS`, and `PASS_WARN_AGE` from /etc/login.defs when adding new entries to /etc/shadow.

`pwunconv` and `grpunconv` are similar. Passwords in the main file are updated from the shadowed file. Entries which exist in the main file but not in the shadowed file are left alone. Finally, the shadowed file is removed. Some password aging information is lost by `pwunconv`. It will convert what it can.

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.		●	●

Internal Only - General

Controls Version	Control	IG 1	IG 2	IG 3
v7	16.4 Encrypt or Hash all Authentication Credentials Encrypt or hash with a salt all authentication credentials when stored.		●	●

Internal Only - General

7.2.2 Ensure /etc/shadow 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.

References:

1. GRID: NIX-00000363
2. NIST SP 800-53 Rev. 5: IA-5
3. NIST SP 800-53 Revision 5 :: CM-6 b
4. NIST SP 800-53A :: CM-6.1 (iv)

Internal Only - General

Additional Information:

```
Red Hat Enterprise Linux 8 Security Technical Implementation Guide
Version 1, Release: 14 Benchmark Date: 24 APR 2024

Vul ID: V-251706
Rule ID: SV-251706r809342
STIG ID: RHEL-08-010121
Severity: CAT II
```

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.		●	●

Internal Only - General

7.2.3 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 to verify all GIDs in `/etc/passwd` exist in `/etc/group`:

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

{
   a_passwd_group_gid=("$(awk -F: '{print $4}' /etc/passwd | sort -u)")
   a_group_gid=("$(awk -F: '{print $3}' /etc/group | sort -u)")
   a_passwd_group_diff=("$(printf '%s\n' "${a_group_gid[@]}"
"${a_passwd_group_gid[@]}" | sort | uniq -u)")
   while IFS= read -r l_gid; do
      awk -F: '($4 == '"$l_gid"') {print "  - User: \"" $1 "\" has GID: \""
$4 "\" which does not exist in /etc/group" }' /etc/passwd
   done < <(printf '%s\n' "${a_passwd_group_gid[@]}"
"${a_passwd_group_diff[@]}" | sort | uniq -D | uniq)
   unset a_passwd_group_gid; unset a_group_gid; unset a_passwd_group_diff
}
```

Nothing should be returned

Remediation:

Analyze the output of the Audit step above and perform the appropriate action to correct any discrepancies found.

References:

1. GRID: NIX-00000365
2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.1 Establish and Maintain an Inventory of Accounts Establish and maintain an inventory of all accounts managed in the enterprise. The inventory must include both user and administrator accounts. The inventory, at a minimum, should contain the person's name, username, start/stop dates, and department. Validate that all active accounts are authorized, on a recurring schedule at a minimum quarterly, or more frequently.	●	●	●
v7	16.6 Maintain an Inventory of Accounts Maintain an inventory of all accounts organized by authentication system.		●	●

Internal Only - General

7.2.4 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:

```
# awk -F: '($1=="shadow") {print $NF}' /etc/group
# awk -F: '($4 == '"$(getent group shadow | awk -F: '{print $3}' | xargs)"')
{print "  - user: \"" $1 "\" primary group is the shadow group"}' /etc/passwd
```

Remediation:

Run the following command to remove all users from the shadow group

```
# sed -ri 's/(^shadow:[^:]*:[^:]*:)([^:]+$)/\1/' /etc/group
```

Change the primary group of any users with shadow as their primary group.

```
# usermod -g <primary group> <user>
```

References:

1. GRID: NIX-00000372
2. NIST SP 800-53 Rev. 5: IA-5

Internal Only - General

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.	●	●	●

Internal Only - General

7.2.5 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.

Satisfies: SRG-OS-000104-GPOS-00051, SRG-OS-000121-GPOS-00062, SRG-OS-000042-GPOS-00020

Audit:

Run the following script and verify no results are returned:

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

{
  while read -r l_count l_uid; do
    if [ "$l_count" -gt 1 ]; then
      echo -e "Duplicate UID: \"$l_uid\" Users: \"$(awk -F: '($3 == n) {
print $1 }' n=$l_uid /etc/passwd | xargs)\""
    fi
  done < <(cut -f3 -d":" /etc/passwd | sort -n | uniq -c)
}
```

Remediation:

Based on the results of the audit script, establish unique UIDs and review all files owned by the shared UIDs to determine which UID they are supposed to belong to.

References:

1. GRID: NIX-00000370
2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5
3. NIST SP 800-53A :: IA-2.1
4. SRG-OS-000104-GPOS-00051
5. SRG-OS-000121-GPOS-00062
6. SRG-OS-000042-GPOS-00020

Internal Only - General

Additional Information:

Red Hat Enterprise Linux 8 Security Technical Implementation Guide Version 1, Release: 14 Benchmark Date: 24 APR 2024

Vul ID: V-230371 Rule ID: SV-230371r627750 STIG ID: RHEL-08-020240 Severity: CAT II

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

7.2.6 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:

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

{
   while read -r l_count l_gid; do
      if [ "$l_count" -gt 1 ]; then
      echo -e "Duplicate GID: \"$l_gid\" Groups: \"$(awk -F: '($3 == n) {
print $1 }' n=$l_gid /etc/group | xargs)\""
    fi
  done < <(cut -f3 -d":" /etc/group | sort -n | uniq -c)
}
```

Remediation:

Based on the results of the audit script, establish unique GIDs and review all files owned by the shared GID to determine which group they are supposed to belong to.

References:

1. GRID: NIX-00000368
2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

Additional Information:

You can also use the `grpck` command to check for other inconsistencies in the `/etc/group` file.

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

7.2.7 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:

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

{
   while read -r l_count l_user; do
      if [ "$l_count" -gt 1 ]; then
         echo -e "Duplicate User: \"$l_user\" Users: \"$(awk -F: '($1 == n) {
print $1 }' n=$l_user /etc/passwd | xargs)\""
      fi
   done < <(cut -f1 -d":" /etc/passwd | sort -n | uniq -c)
}
```

Remediation:

Based on the results of the audit script, establish unique user names for the users. File ownerships will automatically reflect the change as long as the users have unique UIDs.

References:

1. GRID: NIX-00000371
2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

7.2.8 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:

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

{
   while read -r l_count l_group; do
      if [ "$l_count" -gt 1 ]; then
         echo -e "Duplicate Group: \"$l_group\" Groups: \"$(awk -F: '($1 ==
n) { print $1 }' n=$l_group /etc/group | xargs)\""
      fi
   done < <(cut -f1 -d":" /etc/group | sort -n | uniq -c)
}
```

Remediation:

Based on the results of the audit script, establish unique names for the user groups. File group ownerships will automatically reflect the change as long as the groups have unique GIDs.

References:

1. GRID: NIX-00000369
2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

Internal Only - General

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

Internal Only - General

7.2.9 Ensure local interactive user home directories are configured (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. While the system administrator can establish secure permissions for users' home directories, the users can easily override these. Users can be defined in `/etc/passwd` without a home directory or with a home directory that does not actually exist.

Rationale:

Since the user is accountable for files stored in the user home directory, the user must be the owner of the directory. Group or world-writable user home directories may enable malicious users to steal or modify other users' data or to gain another user's system privileges. If the user's home directory does not exist or is unassigned, the user will be placed in "/" and will not be able to write any files or have local environment variables set.

Audit:

Run the following script to ensure:

- local interactive user home directories exist
- local interactive users own their home directories
- local interactive user home directories are mode 750 or more restrictive

Internal Only - General

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

{
   a_output=() a_output2=() a_exists2=() a_mode2=() a_owner2=()
   l_valid_shells="^($( awk -F\/ '$NF != "nologin" {print}' /etc/shells | sed
-rn '/^\//{s,/,\\\\/,g;p}' | paste -s -d '|' - ))$"
   l_mask='0027'; l_max="$( printf '%o' $(( 0777 & ~$l_mask)) )"
   l_users="$(awk -v pat="$l_valid_shells" -F: '$(NF) ~ pat { print $1 " "
$(NF-1) }' /etc/passwd | wc -l)"
   [ "$l_users" -gt 10000 ] && printf '%s\n' "" "  ** INFO **" \
   "  $l_users Local interactive users found on the system" "  This may be a
long running check" "  **********"
   while IFS=" " read -r l_user l_home; do
      if [ -d "$l_home" ]; then
         while IFS=: read -r l_own l_mode; do
            [ "$l_user" != "$l_own" ] && a_owner2+=("  - User: \"$l_user\"
Home \"$l_home\" is owned by: \"$l_own\"")
            [ $(( $l_mode & $l_mask )) -gt 0 ] && a_mode2+=("  - User:
\"$l_user\" Home \"$l_home\" is mode: \"$l_mode\"" \
            "    should be mode: \"$l_max\" or more restrictive")
         done <<< "$(stat -Lc '%U:%#a' "$l_home")"
      else
         a_exists2+=("  - User: \"$l_user\" Home Directory: \"$l_home\"
Doesn't exist")
      fi
   done <<< "$(awk -v pat="$l_valid_shells" -F: '$(NF) ~ pat { print $1 " "
$(NF-1) }' /etc/passwd)"
   [ "${#a_exists2[@]}" -gt 0 ] && a_output2+=("${a_exists2[@]}") || \
   a_output+=("  - All interactive users home directories exist")
   [ "${#a_mode2[@]}" -gt 0 ] && a_output2+=("${a_mode2[@]}") || \
   a_output+=("  - All interactive users home directories are mode \"$l_max\"
or more restrictive")
   [ "${#a_owner2[@]}" -gt 0 ] && a_output2+=("${a_owner2[@]}") || \
   a_output+=("  - All interactive users own their home directory")
   if [ "${#a_output2[@]}" -le 0 ]; then
      printf '%s\n' "" "- Audit Result:" "  ** PASS **" "${a_output[@]}"
   else
      printf '%s\n' "" "- Audit Result:" "  ** FAIL **" " - Reason(s) for
audit failure:" "${a_output2[@]}"
      [ "${#a_output[@]}" -gt 0 ] && printf '%s\n' "- Correctly set:"
"${a_output[@]}"
   fi
}
```

Remediation:

If a local interactive users' home directory is undefined and/or doesn't exist, follow local site policy and perform one of the following:

- Lock the user account
- Remove the user from the system
- Create a directory for the user. If undefined, edit `/etc/passwd` and add the absolute path to the directory to the last field of the user.

Internal Only - General

Run the following script to:

- Remove excessive permissions from local interactive users home directories
- Update the home directory's owner

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

{
   a_output=() a_output2=() a_exists2=() a_mode2=() a_owner2=()
   l_valid_shells="^($( awk -F\/ '$NF != "nologin" {print}' /etc/shells | sed
-rn '/^\//{s,/,\\\\/,g;p}' | paste -s -d '|' - ))$"
   l_mask='0027'; l_max="$( printf '%o' $(( 0777 & ~$l_mask)) )"
   l_users="$(awk -v pat="$l_valid_shells" -F: '$(NF) ~ pat { print $1 " "
$(NF-1) }' /etc/passwd | wc -l)"
   [ "$l_users" -gt 10000 ] && printf '%s\n' "" "  ** INFO **" \
   "  $l_users Local interactive users found on the system" "  This may be a
long running process" "  **********"
   while IFS=" " read -r l_user l_home; do
      if [ -d "$l_home" ]; then
         while IFS=: read -r l_own l_mode; do
            if [ "$l_user" != "$l_own" ]; then
               a_owner2+=("  - User: \"$l_user\" Home \"$l_home\" is owned
by: \"$l_own\"" \
               "    changing owner to: \"$l_user\"") && chown "$l_user"
"$l_home"
            fi
            if [ $(( $l_mode & $l_mask )) -gt 0 ]; then
               a_mode2+=("  - User: \"$l_user\" Home \"$l_home\" is mode:
\"$l_mode\"" \
               "    changing to mode: \"$l_max\" or more restrictive")
               chmod g-w,o-rwx "$l_home"
            fi
         done <<< "$(stat -Lc '%U:%#a' "$l_home")"
      else
         a_exists2+=("  - User: \"$l_user\" Home Directory: \"$l_home\"
Doesn't exist")
      fi
   done <<< "$(awk -v pat="$l_valid_shells" -F: '$(NF) ~ pat { print $1 " "
$(NF-1) }' /etc/passwd)"
   [ "${#a_exists2[@]}" -gt 0 ] && a_output2+=("${a_exists2[@]}")
   [ "${#a_mode2[@]}" -gt 0 ] && a_output2+=("${a_mode2[@]}")
   [ "${#a_owner2[@]}" -gt 0 ] && a_output2+=("${a_owner2[@]}")
   if [ "${#a_output2[@]}" -gt 0 ]; then
      printf '%s\n' "" "${a_output2[@]}"
   else
      printf '%s\n' "" "- No changes required"
   fi
}
```

References:

1. GRID: NIX-00000367
2. NIST SP 800-53 Revision 5 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)

Internal Only - General

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.	●	●	●

Internal Only - General

7.2.10 Ensure local interactive user dot files access is configured (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.

- `.forward` file specifies an email address to forward the user's mail to.
- `.rhosts` file provides the "remote authentication" database for the rcp, rlogin, and rsh commands and the rcmd() function. These files bypass the standard password-based user authentication mechanism. They specify remote hosts and users that are considered trusted (i.e. are allowed to access the local system without supplying a password)
- `.netrc` file contains data for logging into a remote host or passing authentication to an API.
- `.bash_history` file keeps track of the user's commands.

Rationale:

User configuration files with excessive or incorrect access may enable malicious users to steal or modify other users' data or to gain another user's system privileges.

Audit:

Run the following script to verify local interactive user dot files:

- Don't include `.forward`, `.rhosts`, or `.netrc` files
- Are mode `0644` or more restrictive
- Are owned by the local interactive user
- Are group owned by the user's primary group
- `.bash_history` is mode `0600` or more restrictive

Note: If a `.netrc` file is required, and follows local site policy, it should be mode `0600` or more restrictive.

> Note: Home directories located on network/remote filesystems (NFS, CIFS/SMB, sshfs, etc.) are out of scope for this local check and are skipped; audit them from the file server that owns the storage.

Internal Only - General

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

{
 a_output2=(); a_output3=()
 l_maxsize="1000" # Maximum number of local interactive users before warning
(Default 1,000)
 l_valid_shells="^($( awk -F\/ '$NF != "nologin" {print}' /etc/shells | sed -
rn '/^\//{s,/,\\\\/,g;p}' | paste -s -d '|' - ))$"
 a_user_and_home=() # Create array with local users and their home
directories
 while read -r l_local_user l_local_user_home; do # Populate array with users
and user home location
 [[ -n "$l_local_user" && -n "$l_local_user_home" ]] &&
a_user_and_home+=("$l_local_user:$l_local_user_home")
 done <<< "$(awk -v pat="$l_valid_shells" -F: '$(NF) ~ pat { print $1 " "
$(NF-1) }' /etc/passwd)"
 l_asize="${#a_user_and_home[@]}" # Here if we want to look at number of
users before proceeding
 [ "${#a_user_and_home[@]}" -gt "$l_maxsize" ] && printf '%s\n' "" " ** INFO
**" \
 " - \"$l_asize\" Local interactive users found on the system" \
 " - This may be a long running check" ""
 file_access_chk()
 {
 a_access_out=()
 l_max="$( printf '%o' $(( 0777 & ~$l_mask)) )"
 if [ $(( $l_mode & $l_mask )) -gt 0 ]; then
 a_access_out+=(" - File: \"$l_hdfile\" is mode: \"$l_mode\" and should be
mode: \"$l_max\" or more restrictive")
 fi
 if [[ ! "$l_owner" =~ ($l_user) ]]; then
 a_access_out+=(" - File: \"$l_hdfile\" owned by: \"$l_owner\" and should be
owned by \"${l_user//|/ or }\"")
 fi
 if [[ ! "$l_gowner" =~ ($l_group) ]]; then
 a_access_out+=(" - File: \"$l_hdfile\" group owned by: \"$l_gowner\" and
should be group owned by \"${l_group//|/ or }\"")
 fi
 }
 while IFS=: read -r l_user l_home; do
 a_dot_file=(); a_netrc=(); a_netrc_warn=(); a_bhout=(); a_hdirout=()
 # Skip home directories that reside on a network / remote filesystem
 l_fstype="$(findmnt -no FSTYPE --target "$l_home" 2>/dev/null)"
 case "$l_fstype" in
 nfs|nfs4|cifs|smbfs|smb3|fuse.sshfs|afs|ncpfs|glusterfs|ceph)
 a_output3+=(" - User: \"$l_user\" Home Directory: \"$l_home\" skipped
(remote/network filesystem: $l_fstype)")
 continue ;;
 esac
 if [ -d "$l_home" ]; then
 l_group="$(id -gn "$l_user" | xargs)";l_group="${l_group// /|}"
 while IFS= read -r -d $'\0' l_hdfile; do
 while read -r l_mode l_owner l_gowner; do
 case "$(basename "$l_hdfile")" in
 .forward | .rhosts )
 a_dot_file+=(" - File: \"$l_hdfile\" exists") ;;
 .netrc )
```


Internal Only - General

```
 l_mask='0177'; file_access_chk
 if [ "${#a_access_out[@]}" -gt 0 ]; then
 a_netrc+=("${a_access_out[@]}")
 else
 a_netrc_warn+=(" - File: \"$l_hdfile\" exists")
 fi ;;
 .bash_history )
 l_mask='0177'; file_access_chk
 [ "${#a_access_out[@]}" -gt 0 ] && a_bhout+=("${a_access_out[@]}") ;;
 * )
 l_mask='0133'; file_access_chk
 [ "${#a_access_out[@]}" -gt 0 ] && a_hdirout+=("${a_access_out[@]}") ;;
 esac
 done < <(stat -Lc '%#a %U %G' "$l_hdfile")
 done < <(find "$l_home" -xdev -type f -name '.*' -print0)
 fi
 if [[ "${#a_dot_file[@]}" -gt 0 || "${#a_netrc[@]}" -gt 0 ||
"${#a_bhout[@]}" -gt 0 || "${#a_hdirout[@]}" -gt 0 ]]; then
 a_output2+=(" - User: \"$l_user\" Home Directory: \"$l_home\""
"${a_dot_file[@]}" "${a_netrc[@]}" "${a_bhout[@]}" "${a_hdirout[@]}")
 fi
 [ "${#a_netrc_warn[@]}" -gt 0 ] && a_output3+=(" - User: \"$l_user\" Home
Directory: \"$l_home\"" "${a_netrc_warn[@]}")
 done <<< "$(printf '%s\n' "${a_user_and_home[@]}")"
 if [ "${#a_output2[@]}" -le 0 ]; then # If l_output2 is empty, we pass
 [ "${#a_output3[@]}" -gt 0 ] && printf '%s\n' " ** WARNING **"
"${a_output3[@]}"
 printf '%s\n' "- Audit Result:" " ** PASS **"
 else
 printf '%s\n' "- Audit Result:" " ** FAIL **" " - * Reasons for audit
failure * :" "${a_output2[@]}" ""
 [ "${#a_output3[@]}" -gt 0 ] && printf '%s\n' " ** WARNING **"
"${a_output3[@]}"
 fi
}
```

Remediation:

Making global modifications to users' files without alerting the user community can result in unexpected outages and unhappy users. Therefore, it is recommended that a monitoring policy be established to report user dot file permissions and determine the action to be taken in accordance with site policy.

The following script will:

- remove excessive permissions on `dot` files within interactive users' home directories
- change ownership of `dot` files within interactive users' home directories to the user
- change group ownership of `dot` files within interactive users' home directories to the user's primary group
- list `.forward` and `.rhost` files to be investigated and manually deleted

Internal Only - General

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

{
   a_output2=(); a_output3=()
   l_maxsize="1000" # Maximum number of local interactive users before
warning (Default 1,000)
   l_valid_shells="^($( awk -F\/ '$NF != "nologin" {print}' /etc/shells | sed
-rn '/^\//{s,/,\\\\/,g;p}' | paste -s -d '|' - ))$"
   a_user_and_home=() # Create array with local users and their home
directories
   while read -r l_local_user l_local_user_home; do # Populate array with
users and user home location
      [[ -n "$l_local_user" && -n "$l_local_user_home" ]] &&
a_user_and_home+=("$l_local_user:$l_local_user_home")
   done <<< "$(awk -v pat="$l_valid_shells" -F: '$(NF) ~ pat { print $1 " "
$(NF-1) }' /etc/passwd)"
   l_asize="${#a_user_and_home[@]}" # Here if we want to look at number of
users before proceeding
   [ "${#a_user_and_home[@]}" -gt "$l_maxsize" ] && printf '%s\n' "" "  **
INFO **" \
   "  - \"$l_asize\" Local interactive users found on the system" \
   "  - This may be a long running check" ""
   file_access_fix()
   {
      a_access_out=()
      l_max="$( printf '%o' $(( 0777 & ~$l_mask)) )"
      if [ $(( $l_mode & $l_mask )) -gt 0 ]; then
         printf '%s\n' "" "  - File: \"$l_hdfile\" is mode: \"$l_mode\" and
should be mode: \"$l_max\" or more restrictive" \
         "     Updating file: \"$l_hdfile\" to be mode: \"$l_max\" or more
restrictive"
         chmod "$l_change" "$l_hdfile"
      fi
      if [[ ! "$l_owner" =~ ($l_user) ]]; then
         printf '%s\n' "" "  - File: \"$l_hdfile\" owned by: \"$l_owner\" and
should be owned by \"${l_user//|/ or }\"" \
         "     Updating file: \"$l_hdfile\" to be owned by \"${l_user//|/ or
}\""
         chown "$l_user" "$l_hdfile"
      fi
      if [[ ! "$l_gowner" =~ ($l_group) ]]; then
         printf '%s\n' "" "  - File: \"$l_hdfile\" group owned by:
\"$l_gowner\" and should be group owned by \"${l_group//|/ or }\"" \
         "     Updating file: \"$l_hdfile\" to be group owned by
\"${l_group//|/ or }\""
         chgrp "$l_group" "$l_hdfile"
      fi
   }
   while IFS=: read -r l_user l_home; do
      a_dot_file=(); a_netrc=(); a_netrc_warn=(); a_bhout=(); a_hdirout=()
      # Skip home directories that reside on a network / remote filesystem
      l_fstype="$(findmnt -no FSTYPE --target "$l_home" 2>/dev/null)"
      case "$l_fstype" in
         nfs|nfs4|cifs|smbfs|smb3|fuse.sshfs|afs|ncpfs|glusterfs|ceph)
            a_output3+=("   - User: \"$l_user\" Home Directory: \"$l_home\"
skipped (remote/network filesystem: $l_fstype)")
            continue ;;
```


Internal Only - General

```
      esac
      if [ -d "$l_home" ]; then
         l_group="$(id -gn "$l_user" | xargs)";l_group="${l_group// /|}"
         while IFS= read -r -d $'\0' l_hdfile; do
            while read -r l_mode l_owner l_gowner; do
               case "$(basename "$l_hdfile")" in
                  .forward | .rhost )
                     a_dot_file+=("  - File: \"$l_hdfile\" exists" "
Please review and manually delete this file") ;;
                  .netrc )
                     l_mask='0177'; l_change="u-x,go-rwx"; file_access_fix
                     a_netrc_warn+=("   - File: \"$l_hdfile\" exists") ;;
                  .bash_history )
                     l_mask='0177'; l_change="u-x,go-rwx"; file_access_fix ;;
                  * )
                     l_mask='0133'; l_change="u-x,go-wx"; file_access_fix ;;
               esac
            done < <(stat -Lc '%#a %U %G' "$l_hdfile")
         done < <(find "$l_home" -xdev -type f -name '.*' -print0)
      fi
      [ "${#a_dot_file[@]}" -gt 0 ] && a_output2+=(" - User: \"$l_user\" Home
Directory: \"$l_home\"" "${a_dot_file[@]}")
      [ "${#a_netrc_warn[@]}" -gt 0 ] && a_output3+=(" - User: \"$l_user\"
Home Directory: \"$l_home\"" "${a_netrc_warn[@]}")
   done <<< "$(printf '%s\n' "${a_user_and_home[@]}")"
   [ "${#a_output3[@]}" -gt 0 ] && printf '%s\n' "" "  ** WARNING **"
"${a_output3[@]}" ""
   [ "${#a_output2[@]}" -gt 0 ] && printf '%s\n' "" "${a_output2[@]}"
}
```

References:

1. GRID: NIX-00000366
2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	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.	●	●	●

Internal Only - General

7.2.11 Ensure local interactive user dot directories access is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

While the system administrator can establish secure permissions for users' "dot" directories, the users can easily override these. Dot directories are the hidden subdirectories within a user's home directory (their names begin with a `.`) that hold per-user configuration, state, and credential material. Common examples include:

- `.ssh` — holds the user's SSH private keys, `authorized_keys`, `known_hosts`, and

  client configuration. Read or write access by other users can lead to credential theft or unauthorized authentication.

- `.gnupg` — holds the user's GnuPG keyring and trust database.
- `.config`, `.local` — hold application configuration and per-user data

  (XDG base directories).

- `.cache` — holds cached application data, which can include tokens or session material.

Rationale:

User configuration directories with excessive or incorrect access may enable malicious users to read, steal, or modify other users' data — or to gain another user's system privileges. A weakly-permissioned `.ssh` directory is particularly dangerous: if it is group- or world-writable, another user can drop an `authorized_keys` entry and obtain the victim's identity; if it is readable by others, private key material can be exfiltrated.

Impact:

Making global modifications to users' directories without alerting the user community can result in unexpected outages and unhappy users. Tightening permissions on a dot directory may break applications that relied on group or world access to per-user state (for example, a service account configured to read another user's `.config`). Establish a monitoring policy and apply changes in accordance with site policy.

Internal Only - General

This recommendation **intentionally scopes its scan to local interactive users' home directories only** and does not traverse the entire filesystem. As a result, dot directories that live outside those home directories are out of scope for this automated check and must be reviewed separately.

Audit:

Run the following script to verify dot directories within local interactive users' home directories:

- Are mode `0750` or more restrictive
- `.ssh` is mode `0700` or more restrictive
- Are owned by the local interactive user
- Are group owned by the user's primary group

  Note: This audit scans only the home directories of local interactive users (users with a valid login shell, enumerated from `/etc/passwd` and `/etc/shells`). It does **not** perform a full-filesystem traversal.

  Note: Home directories located on network/remote filesystems (NFS, CIFS/SMB, sshfs, etc.) are out of scope for this local check and are skipped; audit them from the file server that owns the storage.

Internal Only - General

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

{
   a_output2=()
   l_maxsize="1000" # Maximum number of local interactive users before
warning (Default 1,000)
   l_valid_shells="^($( awk -F\/ '$NF != "nologin" {print}' /etc/shells | sed
-rn '/^\//{s,/,\\\\/,g;p}' | paste -s -d '|' - ))$"
   a_user_and_home=() # Create array with local users and their home
directories
   while read -r l_local_user l_local_user_home; do # Populate array with
users and user home location
      [[ -n "$l_local_user" && -n "$l_local_user_home" ]] &&
a_user_and_home+=("$l_local_user:$l_local_user_home")
   done <<< "$(awk -v pat="$l_valid_shells" -F: '$(NF) ~ pat { print $1 " "
$(NF-1) }' /etc/passwd)"
   l_asize="${#a_user_and_home[@]}" # Here if we want to look at number of
users before proceeding
   [ "${#a_user_and_home[@]}" -gt "$l_maxsize" ] && printf '%s\n' "" "  **
INFO **" \
   "  - \"$l_asize\" Local interactive users found on the system" \
   "  - This may be a long running check" ""
   dir_access_chk()
   {
      a_access_out=()
      l_max="$( printf '%o' $(( 0777 & ~$l_mask)) )"
      if [ $(( $l_mode & $l_mask )) -gt 0 ]; then
         a_access_out+=("  - Directory: \"$l_hddir\" is mode: \"$l_mode\" and
should be mode: \"$l_max\" or more restrictive")
      fi
      if [[ ! "$l_owner" =~ ($l_user) ]]; then
         a_access_out+=("  - Directory: \"$l_hddir\" owned by: \"$l_owner\"
and should be owned by \"${l_user//|/ or }\"")
      fi
      if [[ ! "$l_gowner" =~ ($l_group) ]]; then
         a_access_out+=("  - Directory: \"$l_hddir\" group owned by:
\"$l_gowner\" and should be group owned by \"${l_group//|/ or }\"")
      fi
   }
   while IFS=: read -r l_user l_home; do
      a_hddirout=()
      # Skip home directories that reside on a network / remote filesystem
      l_fstype="$(findmnt -no FSTYPE --target "$l_home" 2>/dev/null)"
      case "$l_fstype" in
         nfs|nfs4|cifs|smbfs|smb3|fuse.sshfs|afs|ncpfs|glusterfs|ceph)
            printf '%s\n' "" "  - User: \"$l_user\" Home Directory:
\"$l_home\" skipped (remote/network filesystem: $l_fstype)"
            continue ;;
      esac
      if [ -d "$l_home" ]; then
         l_group="$(id -gn "$l_user" | xargs)";l_group="${l_group// /|}"
         while IFS= read -r -d $'\0' l_hddir; do
            while read -r l_mode l_owner l_gowner; do
               case "$(basename "$l_hddir")" in
                  .ssh )
                     l_mask='0077'; dir_access_chk
```


Internal Only - General

```
                    [ "${#a_access_out[@]}" -gt 0 ] &&
a_hddirout+=("${a_access_out[@]}") ;;
                 * )
                    l_mask='0027'; dir_access_chk
                    [ "${#a_access_out[@]}" -gt 0 ] &&
a_hddirout+=("${a_access_out[@]}") ;;
              esac
           done < <(stat -Lc '%#a %U %G' "$l_hddir")
        done < <(find "$l_home" -xdev -type d -name '.*' -print0)
     fi
     if [ "${#a_hddirout[@]}" -gt 0 ]; then
        a_output2+=("  - User: \"$l_user\" Home Directory: \"$l_home\""
"${a_hddirout[@]}")
     fi
  done <<< "$(printf '%s\n' "${a_user_and_home[@]}")"
  if [ "${#a_output2[@]}" -le 0 ]; then # If a_output2 is empty, we pass
     printf '%s\n' "- Audit Result:" "  ** PASS **"
  else
     printf '%s\n' "- Audit Result:" "  ** FAIL **" "  - * Reasons for audit
failure * :" "${a_output2[@]}" ""
  fi
}
```

Remediation:

Making global modifications to users' 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 dot directory permissions and determine the action to be taken in accordance with site policy.

The following script will, within local interactive users' home directories:

- remove excessive permissions on dot directories (max `0750`; `.ssh` max `0700`)
- change ownership of dot directories to the user
- change group ownership of dot directories to the user's primary group

  Note: Like the audit, this remediation scopes its action to local interactive users' home directories only and does not traverse the entire filesystem.

Internal Only - General

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

{
   l_maxsize="1000" # Maximum number of local interactive users before
warning (Default 1,000)
   l_valid_shells="^($( awk -F\/ '$NF != "nologin" {print}' /etc/shells | sed
-rn '/^\//{s,/,\\\\/,g;p}' | paste -s -d '|' - ))$"
   a_user_and_home=() # Create array with local users and their home
directories
   while read -r l_local_user l_local_user_home; do # Populate array with
users and user home location
      [[ -n "$l_local_user" && -n "$l_local_user_home" ]] &&
a_user_and_home+=("$l_local_user:$l_local_user_home")
   done <<< "$(awk -v pat="$l_valid_shells" -F: '$(NF) ~ pat { print $1 " "
$(NF-1) }' /etc/passwd)"
   l_asize="${#a_user_and_home[@]}" # Here if we want to look at number of
users before proceeding
   [ "${#a_user_and_home[@]}" -gt "$l_maxsize" ] && printf '%s\n' "" "  **
INFO **" \
   "  - \"$l_asize\" Local interactive users found on the system" \
   "  - This may be a long running check" ""
   dir_access_fix()
   {
      l_max="$( printf '%o' $(( 0777 & ~$l_mask)) )"
      if [ $(( $l_mode & $l_mask )) -gt 0 ]; then
         printf '%s\n' "" "  - Directory: \"$l_hddir\" is mode: \"$l_mode\"
and should be mode: \"$l_max\" or more restrictive" \
         "     Updating directory: \"$l_hddir\" to be mode: \"$l_max\" or
more restrictive"
         chmod "$l_change" "$l_hddir"
      fi
      if [[ ! "$l_owner" =~ ($l_user) ]]; then
         printf '%s\n' "" "  - Directory: \"$l_hddir\" owned by: \"$l_owner\"
and should be owned by \"${l_user//|/ or }\"" \
         "     Updating directory: \"$l_hddir\" to be owned by \"${l_user//|/
or }\""
         chown "$l_user" "$l_hddir"
      fi
      if [[ ! "$l_gowner" =~ ($l_group) ]]; then
         printf '%s\n' "" "  - Directory: \"$l_hddir\" group owned by:
\"$l_gowner\" and should be group owned by \"${l_group//|/ or }\"" \
         "     Updating directory: \"$l_hddir\" to be group owned by
\"${l_group//|/ or }\""
         chgrp "$l_group" "$l_hddir"
      fi
   }
   while IFS=: read -r l_user l_home; do
      # Skip home directories that reside on a network / remote filesystem
      l_fstype="$(findmnt -no FSTYPE --target "$l_home" 2>/dev/null)"
      case "$l_fstype" in
         nfs|nfs4|cifs|smbfs|smb3|fuse.sshfs|afs|ncpfs|glusterfs|ceph)
            printf '%s\n' "" "  - User: \"$l_user\" Home Directory:
\"$l_home\" skipped (remote/network filesystem: $l_fstype)"
            continue ;;
      esac
      if [ -d "$l_home" ]; then
         l_group="$(id -gn "$l_user" | xargs)";l_group="${l_group// /|}"
```


Internal Only - General

```
         while IFS= read -r -d $'\0' l_hddir; do
            while read -r l_mode l_owner l_gowner; do
               case "$(basename "$l_hddir")" in
                  .ssh )
                     l_mask='0077'; l_change="g-rwx,o-rwx"; dir_access_fix ;;
                  * )
                     l_mask='0027'; l_change="g-w,o-rwx"; dir_access_fix ;;
               esac
            done < <(stat -Lc '%#a %U %G' "$l_hddir")
         done < <(find "$l_home" -xdev -type d -name '.*' -print0)
      fi
   done <<< "$(printf '%s\n' "${a_user_and_home[@]}")"
}
```

Note: This recommendation **intentionally avoids a full-filesystem traversal** (`find /`) when locating dot directories. A system-wide search for `.`-prefixed directories is expensive and can take a very long time on hosts with large or networked filesystems, and it would repeatedly cross into directory trees (caches, build artifacts, application data) that are not user credential stores. Instead, the audit and remediation are **scoped to the home directories of local interactive users** — users who have a valid login shell as enumerated from `/etc/passwd` and `/etc/shells` — and use `find "$l_home" -xdev -type d -name '.*'` (the `-xdev` guard also keeps the scan on the home filesystem and prevents descent into mounted filesystems). This mirrors the scope of recommendation 7.2.10.

Limitation of this scope: Dot directories that hold sensitive material but reside **outside** local interactive users' home directories — for example, credential or configuration directories created under `/opt`, `/srv`, application working directories, or the home directories of service/system accounts that use `nologin`/`false` shells — are **not** evaluated by this automated check. Administrators should ensure, through separate review or site-specific tooling, that no unexpected dot directories with weak permissions (group- or world-writable, or readable by other where they hold secrets) exist outside the audited scope. If your environment stores per-user credentials outside standard interactive home directories, treat those locations under a separate control.

References:

1. NIX-00000409

Internal Only - General

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.	●	●	●

Internal Only - General

Appendix: Summary Table

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1	**Initial Setup**		
1.1	**Filesystem**		
1.1.1	**Configure Filesystem Kernel Modules**		
1.1.1.1	Ensure cramfs kernel module is not available (Automated)	☐	☐
1.1.1.2	Ensure freevxfs kernel module is not available (Automated)	☐	☐
1.1.1.3	Ensure hfs kernel module is not available (Automated)	☐	☐
1.1.1.4	Ensure hfsplus kernel module is not available (Automated)	☐	☐
1.1.1.5	Ensure jffs2 kernel module is not available (Automated)	☐	☐
1.1.1.6	Ensure overlay kernel module is not available (Automated)	☐	☐
1.1.1.7	Ensure squashfs kernel module is not available (Automated)	☐	☐
1.1.1.8	Ensure udf kernel module is not available (Automated)	☐	☐
1.1.1.9	Ensure firewire-core kernel module is not available (Automated)	☐	☐
1.1.1.10	Ensure usb-storage kernel module is not available (Automated)	☐	☐
1.1.1.11	Ensure unused filesystems kernel modules are not available (Manual)	☐	☐
1.1.2	**Configure Filesystem Partitions**		
1.1.2.1	**Configure /tmp**		

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.1.2.1.1	Ensure /tmp is tmpfs or a separate partition (Automated)	☐	☐
1.1.2.1.2	Ensure nodev option set on /tmp partition (Automated)	☐	☐
1.1.2.1.3	Ensure nosuid option set on /tmp partition (Automated)	☐	☐
1.1.2.1.4	Ensure noexec option set on /tmp partition (Automated)	☐	☐
1.1.2.2	**Configure /dev/shm**		
1.1.2.2.1	Ensure /dev/shm is tmpfs or a separate partition (Automated)	☐	☐
1.1.2.2.2	Ensure nodev option set on /dev/shm partition (Automated)	☐	☐
1.1.2.2.3	Ensure nosuid option set on /dev/shm partition (Automated)	☐	☐
1.1.2.2.4	Ensure noexec option set on /dev/shm partition (Automated)	☐	☐
1.1.2.3	**Configure /home**		
1.1.2.3.1	Ensure separate partition exists for /home (Automated)	☐	☐
1.1.2.3.2	Ensure nodev option set on /home partition (Automated)	☐	☐
1.1.2.3.3	Ensure nosuid option set on /home partition (Automated)	☐	☐
1.1.2.4	**Configure /var**		
1.1.2.4.1	Ensure separate partition exists for /var (Automated)	☐	☐
1.1.2.4.2	Ensure nodev option set on /var partition (Automated)	☐	☐
1.1.2.4.3	Ensure nosuid option set on /var partition (Automated)	☐	☐
1.1.2.5	**Configure /var/tmp**		
1.1.2.5.1	Ensure separate partition exists for /var/tmp (Automated)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.1.2.5.2	Ensure nodev option set on /var/tmp partition (Automated)	☐	☐
1.1.2.5.3	Ensure nosuid option set on /var/tmp partition (Automated)	☐	☐
1.1.2.5.4	Ensure noexec option set on /var/tmp partition (Automated)	☐	☐
1.1.2.6	**Configure /var/log**		
1.1.2.6.1	Ensure separate partition exists for /var/log (Automated)	☐	☐
1.1.2.6.2	Ensure nodev option set on /var/log partition (Automated)	☐	☐
1.1.2.6.3	Ensure nosuid option set on /var/log partition (Automated)	☐	☐
1.1.2.6.4	Ensure noexec option set on /var/log partition (Automated)	☐	☐
1.1.2.7	**Configure /var/log/audit**		
1.1.2.7.1	Ensure separate partition exists for /var/log/audit (Automated)	☐	☐
1.1.2.7.2	Ensure nodev option set on /var/log/audit partition (Automated)	☐	☐
1.1.2.7.3	Ensure nosuid option set on /var/log/audit partition (Automated)	☐	☐
1.1.2.7.4	Ensure noexec option set on /var/log/audit partition (Automated)	☐	☐
1.2	**Package Management**		
1.2.1	**Configure Package Repositories**		
1.2.1.1	Ensure the source.list and .source files use the Signed-By option (Manual)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.2.1.2	Ensure weak dependencies are configured (Automated)	☐	☐
1.2.1.3	Ensure access to gpg key files are configured (Automated)	☐	☐
1.2.1.4	Ensure access to /etc/apt/trusted.gpg.d directory is configured (Automated)	☐	☐
1.2.1.5	Ensure access to /etc/apt/auth.conf.d directory is configured (Automated)	☐	☐
1.2.1.6	Ensure access to files in the /etc/apt/auth.conf.d/ directory is configured (Automated)	☐	☐
1.2.1.7	Ensure access to /usr/share/keyrings directory is configured (Automated)	☐	☐
1.2.1.8	Ensure access to /etc/apt/sources.list.d directory is configured (Automated)	☐	☐
1.2.1.9	Ensure access to files in /etc/apt/sources.list.d are configured (Automated)	☐	☐
1.2.1.10	Ensure repositories are set to https in /etc/apt/sources.list (Automated)	☐	☐
1.2.1.11	Ensure repositories are set to https in /etc/apt/sources.list.d directory (Automated)	☐	☐
1.2.1.12	Ensure AllowInsecureRepositories is configured (Automated)	☐	☐
1.2.1.13	Ensure AllowWeakRepositories is configured (Automated)	☐	☐
1.2.1.14	Ensure AllowDowngradeToInsecureRepositories is configured (Automated)	☐	☐
1.2.1.15	Ensure Acquire::Check-Date is configured (Automated)	☐	☐
1.2.2	**Configure Package Updates**		

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.2.2.1	Ensure updates, patches, and additional security software are installed (Manual)	☐	☐
1.3	**Mandatory Access Control**		
1.3.1	**Configure AppArmor**		
1.3.1.1	Ensure apparmor packages are installed (Automated)	☐	☐
1.3.1.2	Ensure AppArmor is enabled (Automated)	☐	☐
1.3.1.3	Ensure all AppArmor Profiles are enforcing (Automated)	☐	☐
1.3.1.4	Ensure kernel.apparmor_restrict_unprivileged_unconfined is enabled (Automated)	☐	☐
1.4	**Configure Bootloader**		
1.4.1	Ensure bootloader password is set (Automated)	☐	☐
1.4.2	Ensure access to bootloader config is configured (Automated)	☐	☐
1.5	**Configure Additional Process Hardening**		
1.5.1	Ensure fs.protected_hardlinks is configured (Automated)	☐	☐
1.5.2	Ensure fs.protected_symlinks is configured (Automated)	☐	☐
1.5.3	Ensure kernel.yama.ptrace_scope is configured (Automated)	☐	☐
1.5.4	Ensure fs.suid_dumpable is configured (Automated)	☐	☐
1.5.5	Ensure kernel.dmesg_restrict is configured (Automated)	☐	☐
1.5.6	Ensure prelink is not installed (Automated)	☐	☐
1.5.7	Ensure Automatic Error Reporting is configured (Automated)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.5.8	Ensure kernel.kptr_restrict is configured (Automated)	☐	☐
1.5.9	Ensure kernel.randomize_va_space is configured (Automated)	☐	☐
1.5.10	Ensure core file size is configured (Automated)	☐	☐
1.5.11	Ensure systemd-coredump ProcessSizeMax is configured (Automated)	☐	☐
1.5.12	Ensure systemd-coredump Storage is configured (Automated)	☐	☐
1.6	**Configure Command Line Warning Banners**		
1.6.1	Ensure /etc/motd is configured (Automated)	☐	☐
1.6.2	Ensure /etc/issue is configured (Automated)	☐	☐
1.6.3	Ensure /etc/issue.net is configured (Automated)	☐	☐
1.6.4	Ensure pam_motd is configured (Automated)	☐	☐
1.6.5	Ensure sshd warning Banner is configured (Automated)	☐	☐
1.6.6	Ensure access to /etc/motd is configured (Automated)	☐	☐
1.6.7	Ensure access to /etc/issue is configured (Automated)	☐	☐
1.6.8	Ensure access to /etc/issue.net is configured (Automated)	☐	☐
1.6.9	Ensure access to pam_motd file is configured (Automated)	☐	☐
1.6.10	Ensure access to sshd warning banner is configured (Automated)	☐	☐
1.6.11	Ensure update-notifier-motd.service is not in use (Automated)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.6.12	Ensure update-notifier-motd.timer is not in use (Automated)	☐	☐
1.7	**Configure GNOME Display Manager**		
1.7.1	Ensure GDM login banner is configured (Automated)	☐	☐
1.7.2	Ensure GDM disable-user-list is configured (Automated)	☐	☐
1.7.3	Ensure GDM screen lock is configured (Automated)	☐	☐
1.7.4	Ensure GDM automount is configured (Automated)	☐	☐
1.7.5	Ensure GDM autorun-never is configured (Automated)	☐	☐
1.7.6	Ensure XDMCP is not enabled (Automated)	☐	☐
1.7.7	Ensure Xwayland is configured (Automated)	☐	☐
2	**Services**		
2.1	**Configure Server Services**		
2.1.1	Ensure autofs services are not in use (Automated)	☐	☐
2.1.2	Ensure avahi daemon services are not in use (Automated)	☐	☐
2.1.3	Ensure dhcp server services are not in use (Automated)	☐	☐
2.1.4	Ensure dns server services are not in use (Automated)	☐	☐
2.1.5	Ensure dnsmasq services are not in use (Automated)	☐	☐
2.1.6	Ensure ftp server services are not in use (Automated)	☐	☐
2.1.7	Ensure ldap server services are not in use (Automated)	☐	☐
2.1.8	Ensure message access server services are not in use (Automated)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
2.1.9	Ensure network file system services are not in use (Automated)	☐	☐
2.1.10	Ensure nis server services are not in use (Automated)	☐	☐
2.1.11	Ensure print server services are not in use (Automated)	☐	☐
2.1.12	Ensure rpcbind services are not in use (Automated)	☐	☐
2.1.13	Ensure rsync services are not in use (Automated)	☐	☐
2.1.14	Ensure samba file server services are not in use (Automated)	☐	☐
2.1.15	Ensure snmp services are not in use (Automated)	☐	☐
2.1.16	Ensure telnet server services are not in use (Automated)	☐	☐
2.1.17	Ensure tftp server services are not in use (Automated)	☐	☐
2.1.18	Ensure web proxy server services are not in use (Automated)	☐	☐
2.1.19	Ensure web server services are not in use (Automated)	☐	☐
2.1.20	Ensure xinetd services are not in use (Automated)	☐	☐
2.1.21	Ensure X window server services are not in use (Automated)	☐	☐
2.1.22	Ensure mail transfer agents are configured for local-only mode (Automated)	☐	☐
2.1.23	Ensure only approved services are listening on a network interface (Manual)	☐	☐
2.2	**Configure Client Services**		
2.2.1	Ensure nis client is not installed (Automated)	☐	☐
2.2.2	Ensure rsh client is not installed (Automated)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
2.2.3	Ensure talk client is not installed (Automated)	☐	☐
2.2.4	Ensure telnet client is not installed (Automated)	☐	☐
2.2.5	Ensure ldap client is not installed (Automated)	☐	☐
2.2.6	Ensure ftp client is not installed (Automated)	☐	☐
2.3	**Configure Time Synchronization**		
2.3.1	**Ensure time synchronization is in use**		
2.3.1.1	Ensure a single time synchronization daemon is in use (Automated)	☐	☐
2.3.2	**Configure systemd-timesyncd**		
2.3.2.1	Ensure systemd-timesyncd configured with authorized timeserver (Automated)	☐	☐
2.3.2.2	Ensure systemd-timesyncd is enabled and running (Automated)	☐	☐
2.3.3	**Configure chrony**		
2.3.3.1	Ensure chrony is configured (Automated)	☐	☐
2.3.3.2	Ensure chrony is running as user _chrony (Automated)	☐	☐
2.3.3.3	Ensure chrony is enabled and running (Automated)	☐	☐
2.4	**Job Schedulers**		
2.4.1	**Configure cron**		
2.4.1.1	Ensure cron daemon is enabled and active (Automated)	☐	☐
2.4.1.2	Ensure access to /etc/crontab is configured (Automated)	☐	☐
2.4.1.3	Ensure access to /etc/cron.hourly is configured (Automated)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
2.4.1.4	Ensure access to /etc/cron.daily is configured (Automated)	☐	☐
2.4.1.5	Ensure access to /etc/cron.weekly is configured (Automated)	☐	☐
2.4.1.6	Ensure access to /etc/cron.monthly is configured (Automated)	☐	☐
2.4.1.7	Ensure access to /etc/cron.yearly is configured (Automated)	☐	☐
2.4.1.8	Ensure access to /etc/cron.d is configured (Automated)	☐	☐
2.4.1.9	Ensure access to crontab is configured (Automated)	☐	☐
2.4.2	**Configure at**		
2.4.2.1	Ensure access to at is configured (Automated)	☐	☐
3	**Network**		
3.1	**Configure Network Devices**		
3.1.1	Ensure IPv6 status is identified (Manual)	☐	☐
3.1.2	Ensure wireless interfaces are not available (Automated)	☐	☐
3.1.3	Ensure bluetooth services are not in use (Automated)	☐	☐
3.2	**Configure Network Kernel Modules**		
3.2.1	Ensure atm kernel module is not available (Automated)	☐	☐
3.2.2	Ensure can kernel module is not available (Automated)	☐	☐
3.2.3	Ensure dccp kernel module is not available (Automated)	☐	☐
3.2.4	Ensure rds kernel module is not available (Automated)	☐	☐
3.2.5	Ensure sctp kernel module is not available (Automated)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
3.2.6	Ensure tipc kernel module is not available (Automated)	☐	☐
3.2.7	Ensure unneeded network protocol kernel modules are not available (Manual)	☐	☐
3.3	**Configure Network Kernel Parameters**		
3.3.1	**Configure IPv4 parameters**		
3.3.1.1	Ensure net.ipv4.ip_forward is configured (Automated)	☐	☐
3.3.1.2	Ensure net.ipv4.conf.all.forwarding is configured (Automated)	☐	☐
3.3.1.3	Ensure net.ipv4.conf.default.forwarding is configured (Automated)	☐	☐
3.3.1.4	Ensure net.ipv4.conf.all.send_redirects is configured (Automated)	☐	☐
3.3.1.5	Ensure net.ipv4.conf.default.send_redirects is configured (Automated)	☐	☐
3.3.1.6	Ensure net.ipv4.icmp_ignore_bogus_error_responses is configured (Automated)	☐	☐
3.3.1.7	Ensure net.ipv4.icmp_echo_ignore_broadcasts is configured (Automated)	☐	☐
3.3.1.8	Ensure net.ipv4.conf.all.accept_redirects is configured (Automated)	☐	☐
3.3.1.9	Ensure net.ipv4.conf.default.accept_redirects is configured (Automated)	☐	☐
3.3.1.10	Ensure net.ipv4.conf.all.secure_redirects is configured (Automated)	☐	☐
3.3.1.11	Ensure net.ipv4.conf.default.secure_redirects is configured (Automated)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
3.3.1.12	Ensure net.ipv4.conf.all.rp_filter is configured (Automated)	☐	☐
3.3.1.13	Ensure net.ipv4.conf.default.rp_filter is configured (Automated)	☐	☐
3.3.1.14	Ensure net.ipv4.conf.all.accept_source_route is configured (Automated)	☐	☐
3.3.1.15	Ensure net.ipv4.conf.default.accept_source_route is configured (Automated)	☐	☐
3.3.1.16	Ensure net.ipv4.conf.all.log_martians is configured (Automated)	☐	☐
3.3.1.17	Ensure net.ipv4.conf.default.log_martians is configured (Automated)	☐	☐
3.3.1.18	Ensure net.ipv4.tcp_syncookies is configured (Automated)	☐	☐
3.3.1.19	Ensure net.ipv4.conf.all.route_localnet is configured (Automated)	☐	☐
3.3.2	**Configure IPv6 parameters**		
3.3.2.1	Ensure net.ipv6.conf.all.forwarding is configured (Automated)	☐	☐
3.3.2.2	Ensure net.ipv6.conf.default.forwarding is configured (Automated)	☐	☐
3.3.2.3	Ensure net.ipv6.conf.all.accept_redirects is configured (Automated)	☐	☐
3.3.2.4	Ensure net.ipv6.conf.default.accept_redirects is configured (Automated)	☐	☐
3.3.2.5	Ensure net.ipv6.conf.all.accept_source_route is configured (Automated)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
3.3.2.6	Ensure net.ipv6.conf.default.accept_source_route is configured (Automated)	☐	☐
3.3.2.7	Ensure net.ipv6.conf.all.accept_ra is configured (Automated)	☐	☐
3.3.2.8	Ensure net.ipv6.conf.default.accept_ra is configured (Automated)	☐	☐
4	**Host Based Firewall**		
4.1	**Configure Uncomplicated Firewall**		
4.1.1	Ensure ufw is installed (Automated)	☐	☐
4.1.2	Ensure ufw service is configured (Automated)	☐	☐
4.1.3	Ensure ufw incoming default is configured (Automated)	☐	☐
4.1.4	Ensure ufw outgoing default is configured (Automated)	☐	☐
4.1.5	Ensure ufw routed default is configured (Automated)	☐	☐
5	**Access Control**		
5.1	**Configure SSH Server**		
5.1.1	Ensure access to /etc/ssh/sshd_config is configured (Automated)	☐	☐
5.1.2	Ensure access to /etc/ssh/sshd_config.d/ directory is configured (Automated)	☐	☐
5.1.3	Ensure access to SSH private host key files is configured (Automated)	☐	☐
5.1.4	Ensure access to SSH public host key files is configured (Automated)	☐	☐
5.1.5	Ensure sshd access is configured (Automated)	☐	☐
5.1.6	Ensure sshd Banner is configured (Automated)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
5.1.7	Ensure sshd Ciphers are configured (Automated)	☐	☐
5.1.8	Ensure sshd ClientAliveInterval and ClientAliveCountMax are configured (Automated)	☐	☐
5.1.9	Ensure sshd DisableForwarding is enabled (Automated)	☐	☐
5.1.10	Ensure sshd GSSAPIAuthentication is disabled (Automated)	☐	☐
5.1.11	Ensure sshd HostbasedAuthentication is disabled (Automated)	☐	☐
5.1.12	Ensure sshd IgnoreRhosts is enabled (Automated)	☐	☐
5.1.13	Ensure sshd KexAlgorithms is configured (Automated)	☐	☐
5.1.14	Ensure sshd post-quantum cryptography key exchange algorithms are configured (Manual)	☐	☐
5.1.15	Ensure sshd LoginGraceTime is configured (Automated)	☐	☐
5.1.16	Ensure sshd LogLevel is configured (Automated)	☐	☐
5.1.17	Ensure sshd MACs are configured (Automated)	☐	☐
5.1.18	Ensure sshd MaxAuthTries is configured (Automated)	☐	☐
5.1.19	Ensure sshd MaxSessions is configured (Automated)	☐	☐
5.1.20	Ensure sshd MaxStartups is configured (Automated)	☐	☐
5.1.21	Ensure sshd PermitEmptyPasswords is disabled (Automated)	☐	☐
5.1.22	Ensure sshd PermitRootLogin is disabled (Automated)	☐	☐
5.1.23	Ensure sshd PermitUserEnvironment is disabled (Automated)	☐	☐
5.1.24	Ensure sshd UsePAM is enabled (Automated)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
5.2	**Configure privilege escalation**		
5.2.1	Ensure sudo is installed (Automated)	☐	☐
5.2.2	Ensure sudo commands use pty (Automated)	☐	☐
5.2.3	Ensure sudo log file exists (Automated)	☐	☐
5.2.4	Ensure users must provide password for escalation (Automated)	☐	☐
5.2.5	Ensure re-authentication for privilege escalation is not disabled globally (Automated)	☐	☐
5.2.6	Ensure sudo timestamp_timeout is configured (Automated)	☐	☐
5.2.7	Ensure access to the su command is restricted (Automated)	☐	☐
5.3	**Pluggable Authentication Modules**		
5.3.1	**Configure PAM software packages**		
5.3.1.1	Ensure latest version of pam is installed (Automated)	☐	☐
5.3.1.2	Ensure latest version of libpam-modules is installed (Automated)	☐	☐
5.3.1.3	Ensure latest version of libpam-pwquality is installed (Automated)	☐	☐
5.3.1.4	Ensure latest version of cracklib-runtime is installed (Automated)	☐	☐
5.3.2	**Configure pam-auth-update profiles**		
5.3.2.1	Ensure pam_unix module is enabled (Automated)	☐	☐
5.3.2.2	Ensure pam_faillock module is enabled (Automated)	☐	☐
5.3.2.3	Ensure pam_pwquality module is enabled (Automated)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
5.3.2.4	Ensure pam_pwhistory module is enabled (Automated)	☐	☐
5.3.3	**Configure PAM Arguments**		
5.3.3.1	**Configure pam_faillock module**		
5.3.3.1.1	Ensure password failed attempts lockout is configured (Automated)	☐	☐
5.3.3.1.2	Ensure password unlock time is configured (Automated)	☐	☐
5.3.3.1.3	Ensure password failed attempts lockout includes root account (Automated)	☐	☐
5.3.3.2	**Configure pam_pwquality module**		
5.3.3.2.1	Ensure password number of changed characters is configured (Automated)	☐	☐
5.3.3.2.2	Ensure password length is configured (Automated)	☐	☐
5.3.3.2.3	Ensure password complexity is configured (Manual)	☐	☐
5.3.3.2.4	Ensure password same consecutive characters is configured (Automated)	☐	☐
5.3.3.2.5	Ensure password maximum sequential characters is configured (Automated)	☐	☐
5.3.3.2.6	Ensure password dictionary check is enabled (Automated)	☐	☐
5.3.3.2.7	Ensure password quality checking is enforced (Automated)	☐	☐
5.3.3.2.8	Ensure password quality is enforced for the root user (Automated)	☐	☐
5.3.3.3	**Configure pam_pwhistory module**		
5.3.3.3.1	Ensure password history remember is configured (Automated)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
5.3.3.3.2	Ensure password history is enforced for the root user (Automated)	☐	☐
5.3.3.3.3	Ensure pam_pwhistory includes use_authtok (Automated)	☐	☐
5.3.3.4	**Configure pam_unix module**		
5.3.3.4.1	Ensure pam_unix does not include nullok (Automated)	☐	☐
5.3.3.4.2	Ensure pam_unix does not include remember (Automated)	☐	☐
5.3.3.4.3	Ensure pam_unix includes a strong password hashing algorithm (Automated)	☐	☐
5.3.3.4.4	Ensure pam_unix includes use_authtok (Automated)	☐	☐
5.4	**User Accounts and Environment**		
5.4.1	**Configure shadow password suite parameters**		
5.4.1.1	Ensure password expiration is configured (Automated)	☐	☐
5.4.1.2	Ensure minimum password days is configured (Manual)	☐	☐
5.4.1.3	Ensure password expiration warning days is configured (Automated)	☐	☐
5.4.1.4	Ensure strong password hashing algorithm is configured (Automated)	☐	☐
5.4.1.5	Ensure inactive password lock is configured (Automated)	☐	☐
5.4.1.6	Ensure all users last password change date is in the past (Automated)	☐	☐
5.4.2	**Configure root and system accounts and environment**		
5.4.2.1	Ensure root is the only UID 0 account (Automated)	☐	☐
5.4.2.2	Ensure root is the only GID 0 account (Automated)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
5.4.2.3	Ensure group root is the only GID 0 group (Automated)	☐	☐
5.4.2.4	Ensure root account access is controlled (Automated)	☐	☐
5.4.2.5	Ensure root path integrity (Automated)	☐	☐
5.4.2.6	Ensure root user umask is configured (Automated)	☐	☐
5.4.2.7	Ensure system accounts do not have a valid login shell (Automated)	☐	☐
5.4.2.8	Ensure accounts without a valid login shell are locked (Automated)	☐	☐
5.4.3	**Configure user default environment**		
5.4.3.1	Ensure nologin is not listed in /etc/shells (Automated)	☐	☐
5.4.3.2	Ensure default user shell timeout is configured (Automated)	☐	☐
5.4.3.3	Ensure default user umask is configured (Automated)	☐	☐
6	**Logging and Auditing**		
6.1	**System Logging**		
6.1.1	**Configure journald**		
6.1.1.1	**Configure journald**		
6.1.1.1.1	Ensure journald service is active (Automated)	☐	☐
6.1.1.1.2	Ensure systemd-journal-remote service is not in use (Automated)	☐	☐
6.1.1.1.3	Ensure journald is configured to send logs to rsyslog (Automated)	☐	☐
6.1.1.1.4	Ensure journald log file access is configured (Manual)	☐	☐
6.1.1.1.5	Ensure journald log file rotation is configured (Manual)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
6.1.1.1.6	Ensure journald Storage is configured (Automated)	☐	☐
6.1.1.1.7	Ensure journald Compress is configured (Automated)	☐	☐
6.1.2	**Configure rsyslog**		
6.1.2.1	Ensure rsyslog is installed (Automated)	☐	☐
6.1.2.2	Ensure rsyslog service is enabled and active (Automated)	☐	☐
6.1.2.3	Ensure rsyslog log file creation mode is configured (Automated)	☐	☐
6.1.2.4	Ensure rsyslog logging is configured (Manual)	☐	☐
6.1.2.5	Ensure rsyslog is configured to send logs to a remote log host (Manual)	☐	☐
6.1.2.6	Ensure rsyslog is not configured to receive logs from a remote client (Automated)	☐	☐
6.1.2.7	Ensure logrotate is configured (Manual)	☐	☐
6.1.2.8	Ensure rsyslog-gnutls is installed (Automated)	☐	☐
6.1.2.9	Ensure rsyslog forwarding uses gtls (Automated)	☐	☐
6.1.2.10	Ensure rsyslog CA certificates are configured (Manual)	☐	☐
6.1.3	**Configure Logfiles**		
6.1.3.1	Ensure access to all logfiles has been configured (Automated)	☐	☐
6.2	**System Auditing**		
6.2.1	**Configure auditd Service**		
6.2.1.1	Ensure auditd packages are installed (Automated)	☐	☐
6.2.1.2	Ensure auditd service is enabled and active (Automated)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
6.2.1.3	Ensure auditing for processes that start prior to auditd is enabled (Automated)	☐	☐
6.2.1.4	Ensure audit_backlog_limit is configured (Automated)	☐	☐
6.2.2	**Configure Data Retention**		
6.2.2.1	Ensure audit log storage size is configured (Automated)	☐	☐
6.2.2.2	Ensure audit logs are not automatically deleted (Automated)	☐	☐
6.2.2.3	Ensure system is disabled when audit logs are full (Automated)	☐	☐
6.2.2.4	Ensure system warns when audit logs are low on space (Automated)	☐	☐
6.2.3	**Configure auditd Rules**		
6.2.3.1	Ensure modification of the /etc/sudoers file is collected (Automated)	☐	☐
6.2.3.2	Ensure actions as another user are always logged (Automated)	☐	☐
6.2.3.3	Ensure events that modify the sudo log file are collected (Automated)	☐	☐
6.2.3.4	Ensure events that modify date and time information are collected (Automated)	☐	☐
6.2.3.5	Ensure events that modify sethostname and setdomainname are collected (Automated)	☐	☐
6.2.3.6	Ensure events that modify /etc/issue and /etc/issue.net are collected (Automated)	☐	☐
6.2.3.7	Ensure events that modify /etc/hosts and /etc/hostname are collected (Automated)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
6.2.3.8	Ensure events that modify the system's network environment are collected (Automated)	☐	☐
6.2.3.9	Ensure events that modify /etc/NetworkManager directory are collected (Automated)	☐	☐
6.2.3.10	Ensure use of privileged commands are collected (Automated)	☐	☐
6.2.3.11	Ensure unsuccessful file access attempts are collected (Automated)	☐	☐
6.2.3.12	Ensure events that modify /etc/group information are collected (Automated)	☐	☐
6.2.3.13	Ensure events that modify /etc/passwd information are collected (Automated)	☐	☐
6.2.3.14	Ensure events that modify /etc/shadow and /etc/gshadow are collected (Automated)	☐	☐
6.2.3.15	Ensure events that modify /etc/security/opasswd are collected (Automated)	☐	☐
6.2.3.16	Ensure events that modify /etc/nsswitch.conf file are collected (Automated)	☐	☐
6.2.3.17	Ensure events that modify /etc/pam.conf and /etc/pam.d/ information are collected (Automated)	☐	☐
6.2.3.18	Ensure discretionary access control permission modification events chmod,fchmod,fchmodat,fchmodat2 are collected (Automated)	☐	☐
6.2.3.19	Ensure discretionary access control permission modification events chown,fchown,lchown,fchownat are collected (Automated)	☐	☐
6.2.3.20	Ensure discretionary access control permission modification events setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr collected (Automated)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
6.2.3.21	Ensure successful file system mounts are collected (Automated)	☐	☐
6.2.3.22	Ensure session initiation information is collected (Automated)	☐	☐
6.2.3.23	Ensure login and logout events are collected (Automated)	☐	☐
6.2.3.24	Ensure unlink file deletion events by users are collected (Automated)	☐	☐
6.2.3.25	Ensure rename file deletion events by users are collected (Automated)	☐	☐
6.2.3.26	Ensure events that modify the system's Mandatory Access Controls are collected (Automated)	☐	☐
6.2.3.27	Ensure successful and unsuccessful attempts to use the chcon command are collected (Automated)	☐	☐
6.2.3.28	Ensure successful and unsuccessful attempts to use the setfacl command are collected (Automated)	☐	☐
6.2.3.29	Ensure successful and unsuccessful attempts to use the chacl command are collected (Automated)	☐	☐
6.2.3.30	Ensure successful and unsuccessful attempts to use the usermod command are collected (Automated)	☐	☐
6.2.3.31	Ensure kernel module loading unloading and modification is collected (Automated)	☐	☐
6.2.3.32	Ensure kernel "init_module" and "finit_module" loading unloading and modification is collected (Automated)	☐	☐
6.2.3.33	Ensure kernel "delete_module" loading unloading and modification is collected (Automated)	☐	☐
6.2.3.34	Ensure the audit configuration is loaded regardless of errors (Automated)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
6.2.3.35	Ensure the audit configuration is immutable (Automated)	☐	☐
6.2.3.36	Ensure the running and on disk configuration is the same (Manual)	☐	☐
6.2.4	**Configure auditd File Access**		
6.2.4.1	Ensure audit log files mode is configured (Automated)	☐	☐
6.2.4.2	Ensure audit log files owner is configured (Automated)	☐	☐
6.2.4.3	Ensure audit log files group owner is configured (Automated)	☐	☐
6.2.4.4	Ensure the audit log file directory mode is configured (Automated)	☐	☐
6.2.4.5	Ensure audit configuration files mode is configured (Automated)	☐	☐
6.2.4.6	Ensure audit configuration files owner is configured (Automated)	☐	☐
6.2.4.7	Ensure audit configuration files group owner is configured (Automated)	☐	☐
6.2.4.8	Ensure audit tools mode is configured (Automated)	☐	☐
6.2.4.9	Ensure audit tools owner is configured (Automated)	☐	☐
6.2.4.10	Ensure audit tools group owner is configured (Automated)	☐	☐
6.3	**Configure Integrity Checking**		
6.3.1	Ensure AIDE is installed (Automated)	☐	☐
6.3.2	Ensure filesystem integrity is regularly checked (Automated)	☐	☐
6.3.3	Ensure cryptographic mechanisms are used to protect the integrity of audit tools (Automated)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
7	**System Maintenance**		
7.1	**Configure system file and directory access**		
7.1.1	Ensure access to /etc/passwd is configured (Automated)	☐	☐
7.1.2	Ensure access to /etc/passwd- is configured (Automated)	☐	☐
7.1.3	Ensure access to /etc/group is configured (Automated)	☐	☐
7.1.4	Ensure access to /etc/group- is configured (Automated)	☐	☐
7.1.5	Ensure access to /etc/shadow is configured (Automated)	☐	☐
7.1.6	Ensure access to /etc/shadow- is configured (Automated)	☐	☐
7.1.7	Ensure access to /etc/gshadow is configured (Automated)	☐	☐
7.1.8	Ensure access to /etc/gshadow- is configured (Automated)	☐	☐
7.1.9	Ensure access to /etc/shells is configured (Automated)	☐	☐
7.1.10	Ensure access to /etc/security/opasswd is configured (Automated)	☐	☐
7.1.11	Ensure world writable files and directories are secured (Automated)	☐	☐
7.1.12	Ensure no files or directories without an owner and a group exist (Automated)	☐	☐
7.1.13	Ensure SUID and SGID files are reviewed (Manual)	☐	☐
7.2	**Local User and Group Settings**		
7.2.1	Ensure accounts in /etc/passwd use shadowed passwords (Automated)	☐	☐

Internal Only - General

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
7.2.2	Ensure /etc/shadow password fields are not empty (Automated)	☐	☐
7.2.3	Ensure all groups in /etc/passwd exist in /etc/group (Automated)	☐	☐
7.2.4	Ensure shadow group is empty (Automated)	☐	☐
7.2.5	Ensure no duplicate UIDs exist (Automated)	☐	☐
7.2.6	Ensure no duplicate GIDs exist (Automated)	☐	☐
7.2.7	Ensure no duplicate user names exist (Automated)	☐	☐
7.2.8	Ensure no duplicate group names exist (Automated)	☐	☐
7.2.9	Ensure local interactive user home directories are configured (Automated)	☐	☐
7.2.10	Ensure local interactive user dot files access is configured (Automated)	☐	☐
7.2.11	Ensure local interactive user dot directories access is configured (Automated)	☐	☐

Internal Only - General

Appendix: CIS Controls v7 IG 1 Mapped Recommendations

Recommendation		Set Correctly	
		Yes	No
1.1.2.1.1	Ensure /tmp is tmpfs or a separate partition	☐	☐
1.1.2.1.2	Ensure nodev option set on /tmp partition	☐	☐
1.1.2.1.3	Ensure nosuid option set on /tmp partition	☐	☐
1.1.2.1.4	Ensure noexec option set on /tmp partition	☐	☐
1.1.2.2.1	Ensure /dev/shm is tmpfs or a separate partition	☐	☐
1.1.2.2.2	Ensure nodev option set on /dev/shm partition	☐	☐
1.1.2.2.3	Ensure nosuid option set on /dev/shm partition	☐	☐
1.1.2.2.4	Ensure noexec option set on /dev/shm partition	☐	☐
1.1.2.3.1	Ensure separate partition exists for /home	☐	☐
1.1.2.3.2	Ensure nodev option set on /home partition	☐	☐
1.1.2.3.3	Ensure nosuid option set on /home partition	☐	☐
1.1.2.4.1	Ensure separate partition exists for /var	☐	☐
1.1.2.4.2	Ensure nodev option set on /var partition	☐	☐
1.1.2.4.3	Ensure nosuid option set on /var partition	☐	☐
1.1.2.5.2	Ensure nodev option set on /var/tmp partition	☐	☐
1.1.2.5.3	Ensure nosuid option set on /var/tmp partition	☐	☐
1.1.2.5.4	Ensure noexec option set on /var/tmp partition	☐	☐
1.1.2.6.2	Ensure nodev option set on /var/log partition	☐	☐
1.1.2.6.3	Ensure nosuid option set on /var/log partition	☐	☐
1.1.2.6.4	Ensure noexec option set on /var/log partition	☐	☐
1.1.2.7.2	Ensure nodev option set on /var/log/audit partition	☐	☐
1.1.2.7.3	Ensure nosuid option set on /var/log/audit partition	☐	☐
1.1.2.7.4	Ensure noexec option set on /var/log/audit partition	☐	☐
1.2.1.1	Ensure the source.list and .source files use the Signed-By option	☐	☐
1.2.1.3	Ensure access to gpg key files are configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
1.2.1.4	Ensure access to /etc/apt/trusted.gpg.d directory is configured	☐	☐
1.2.1.5	Ensure access to /etc/apt/auth.conf.d directory is configured	☐	☐
1.2.1.6	Ensure access to files in the /etc/apt/auth.conf.d/ directory is configured	☐	☐
1.2.1.8	Ensure access to /etc/apt/sources.list.d directory is configured	☐	☐
1.2.1.9	Ensure access to files in /etc/apt/sources.list.d are configured	☐	☐
1.2.1.12	Ensure AllowInsecureRepositories is configured	☐	☐
1.2.1.13	Ensure AllowWeakRepositories is configured	☐	☐
1.2.1.14	Ensure AllowDowngradeToInsecureRepositories is configured	☐	☐
1.2.2.1	Ensure updates, patches, and additional security software are installed	☐	☐
1.3.1.1	Ensure apparmor packages are installed	☐	☐
1.3.1.2	Ensure AppArmor is enabled	☐	☐
1.3.1.3	Ensure all AppArmor Profiles are enforcing	☐	☐
1.4.1	Ensure bootloader password is set	☐	☐
1.4.2	Ensure access to bootloader config is configured	☐	☐
1.5.5	Ensure kernel.dmesg_restrict is configured	☐	☐
1.6.6	Ensure access to /etc/motd is configured	☐	☐
1.6.7	Ensure access to /etc/issue is configured	☐	☐
1.6.8	Ensure access to /etc/issue.net is configured	☐	☐
1.6.9	Ensure access to pam_motd file is configured	☐	☐
1.6.10	Ensure access to sshd warning banner is configured	☐	☐
1.7.3	Ensure GDM screen lock is configured	☐	☐
1.7.4	Ensure GDM automount is configured	☐	☐
1.7.5	Ensure GDM autorun-never is configured	☐	☐
2.1.1	Ensure autofs services are not in use	☐	☐
2.1.16	Ensure telnet server services are not in use	☐	☐
2.1.21	Ensure X window server services are not in use	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
2.2.1	Ensure nis client is not installed	☐	☐
2.4.1.2	Ensure access to /etc/crontab is configured	☐	☐
2.4.1.3	Ensure access to /etc/cron.hourly is configured	☐	☐
2.4.1.4	Ensure access to /etc/cron.daily is configured	☐	☐
2.4.1.5	Ensure access to /etc/cron.weekly is configured	☐	☐
2.4.1.6	Ensure access to /etc/cron.monthly is configured	☐	☐
2.4.1.7	Ensure access to /etc/cron.yearly is configured	☐	☐
2.4.1.8	Ensure access to /etc/cron.d is configured	☐	☐
2.4.1.9	Ensure access to crontab is configured	☐	☐
2.4.2.1	Ensure access to at is configured	☐	☐
3.3.1.16	Ensure net.ipv4.conf.all.log_martians is configured	☐	☐
3.3.1.17	Ensure net.ipv4.conf.default.log_martians is configured	☐	☐
4.1.1	Ensure ufw is installed	☐	☐
4.1.2	Ensure ufw service is configured	☐	☐
4.1.3	Ensure ufw incoming default is configured	☐	☐
4.1.4	Ensure ufw outgoing default is configured	☐	☐
4.1.5	Ensure ufw routed default is configured	☐	☐
5.1.1	Ensure access to /etc/ssh/sshd_config is configured	☐	☐
5.1.2	Ensure access to /etc/ssh/sshd_config.d/ directory is configured	☐	☐
5.1.3	Ensure access to SSH private host key files is configured	☐	☐
5.1.4	Ensure access to SSH public host key files is configured	☐	☐
5.1.5	Ensure sshd access is configured	☐	☐
5.1.16	Ensure sshd LogLevel is configured	☐	☐
5.1.22	Ensure sshd PermitRootLogin is disabled	☐	☐
5.2.1	Ensure sudo is installed	☐	☐
5.2.2	Ensure sudo commands use pty	☐	☐
5.2.4	Ensure users must provide password for escalation	☐	☐
5.2.5	Ensure re-authentication for privilege escalation is not disabled globally	☐	☐
5.2.6	Ensure sudo timestamp_timeout is configured	☐	☐
5.2.7	Ensure access to the su command is restricted	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
5.3.2.1	Ensure pam_unix module is enabled	☐	☐
5.4.2.1	Ensure root is the only UID 0 account	☐	☐
5.4.2.2	Ensure root is the only GID 0 account	☐	☐
5.4.2.3	Ensure group root is the only GID 0 group	☐	☐
5.4.2.4	Ensure root account access is controlled	☐	☐
5.4.2.6	Ensure root user umask is configured	☐	☐
5.4.2.7	Ensure system accounts do not have a valid login shell	☐	☐
5.4.2.8	Ensure accounts without a valid login shell are locked	☐	☐
5.4.3.2	Ensure default user shell timeout is configured	☐	☐
5.4.3.3	Ensure default user umask is configured	☐	☐
6.1.1.1.1	Ensure journald service is active	☐	☐
6.1.1.1.3	Ensure journald is configured to send logs to rsyslog	☐	☐
6.1.1.1.4	Ensure journald log file access is configured	☐	☐
6.1.1.1.5	Ensure journald log file rotation is configured	☐	☐
6.1.1.1.6	Ensure journald Storage is configured	☐	☐
6.1.1.1.7	Ensure journald Compress is configured	☐	☐
6.1.2.1	Ensure rsyslog is installed	☐	☐
6.1.2.2	Ensure rsyslog service is enabled and active	☐	☐
6.1.2.3	Ensure rsyslog log file creation mode is configured	☐	☐
6.1.2.4	Ensure rsyslog logging is configured	☐	☐
6.1.2.5	Ensure rsyslog is configured to send logs to a remote log host	☐	☐
6.1.2.8	Ensure rsyslog-gnutls is installed	☐	☐
6.1.3.1	Ensure access to all logfiles has been configured	☐	☐
6.2.1.1	Ensure auditd packages are installed	☐	☐
6.2.1.2	Ensure auditd service is enabled and active	☐	☐
6.2.1.3	Ensure auditing for processes that start prior to auditd is enabled	☐	☐
6.2.1.4	Ensure audit_backlog_limit is configured	☐	☐
6.2.3.10	Ensure use of privileged commands are collected	☐	☐
6.2.3.23	Ensure login and logout events are collected	☐	☐
6.2.3.24	Ensure unlink file deletion events by users are collected	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
6.2.3.25	Ensure rename file deletion events by users are collected	☐	☐
6.2.3.27	Ensure successful and unsuccessful attempts to use the chcon command are collected	☐	☐
6.2.3.28	Ensure successful and unsuccessful attempts to use the setfacl command are collected	☐	☐
6.2.3.29	Ensure successful and unsuccessful attempts to use the chacl command are collected	☐	☐
6.2.3.30	Ensure successful and unsuccessful attempts to use the usermod command are collected	☐	☐
6.2.3.31	Ensure kernel module loading unloading and modification is collected	☐	☐
6.2.3.32	Ensure kernel "init_module" and "finit_module" loading unloading and modification is collected	☐	☐
6.2.3.33	Ensure kernel "delete_module" loading unloading and modification is collected	☐	☐
6.2.3.34	Ensure the audit configuration is loaded regardless of errors	☐	☐
6.2.3.35	Ensure the audit configuration is immutable	☐	☐
6.2.4.1	Ensure audit log files mode is configured	☐	☐
6.2.4.2	Ensure audit log files owner is configured	☐	☐
6.2.4.3	Ensure audit log files group owner is configured	☐	☐
6.2.4.4	Ensure the audit log file directory mode is configured	☐	☐
6.2.4.5	Ensure audit configuration files mode is configured	☐	☐
6.2.4.6	Ensure audit configuration files owner is configured	☐	☐
6.2.4.7	Ensure audit configuration files group owner is configured	☐	☐
6.2.4.8	Ensure audit tools mode is configured	☐	☐
6.2.4.9	Ensure audit tools owner is configured	☐	☐
6.2.4.10	Ensure audit tools group owner is configured	☐	☐
7.1.1	Ensure access to /etc/passwd is configured	☐	☐
7.1.2	Ensure access to /etc/passwd- is configured	☐	☐
7.1.3	Ensure access to /etc/group is configured	☐	☐
7.1.4	Ensure access to /etc/group- is configured	☐	☐
7.1.5	Ensure access to /etc/shadow is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
7.1.6	Ensure access to /etc/shadow- is configured	☐	☐
7.1.7	Ensure access to /etc/gshadow is configured	☐	☐
7.1.8	Ensure access to /etc/gshadow- is configured	☐	☐
7.1.9	Ensure access to /etc/shells is configured	☐	☐
7.1.10	Ensure access to /etc/security/opasswd is configured	☐	☐
7.1.11	Ensure world writable files and directories are secured	☐	☐
7.1.12	Ensure no files or directories without an owner and a group exist	☐	☐
7.1.13	Ensure SUID and SGID files are reviewed	☐	☐
7.2.4	Ensure shadow group is empty	☐	☐
7.2.9	Ensure local interactive user home directories are configured	☐	☐
7.2.10	Ensure local interactive user dot files access is configured	☐	☐
7.2.11	Ensure local interactive user dot directories access is configured	☐	☐

Internal Only - General

Appendix: CIS Controls v7 IG 2 Mapped Recommendations

Recommendation		Set Correctly	
		Yes	No
1.1.1.1	Ensure cramfs kernel module is not available	☐	☐
1.1.1.2	Ensure freevxfs kernel module is not available	☐	☐
1.1.1.3	Ensure hfs kernel module is not available	☐	☐
1.1.1.4	Ensure hfsplus kernel module is not available	☐	☐
1.1.1.5	Ensure jffs2 kernel module is not available	☐	☐
1.1.1.6	Ensure overlay kernel module is not available	☐	☐
1.1.1.7	Ensure squashfs kernel module is not available	☐	☐
1.1.1.8	Ensure udf kernel module is not available	☐	☐
1.1.1.9	Ensure firewire-core kernel module is not available	☐	☐
1.1.1.10	Ensure usb-storage kernel module is not available	☐	☐
1.1.1.11	Ensure unused filesystems kernel modules are not available	☐	☐
1.1.2.1.1	Ensure /tmp is tmpfs or a separate partition	☐	☐
1.1.2.1.2	Ensure nodev option set on /tmp partition	☐	☐
1.1.2.1.3	Ensure nosuid option set on /tmp partition	☐	☐
1.1.2.1.4	Ensure noexec option set on /tmp partition	☐	☐
1.1.2.2.1	Ensure /dev/shm is tmpfs or a separate partition	☐	☐
1.1.2.2.2	Ensure nodev option set on /dev/shm partition	☐	☐
1.1.2.2.3	Ensure nosuid option set on /dev/shm partition	☐	☐
1.1.2.2.4	Ensure noexec option set on /dev/shm partition	☐	☐
1.1.2.3.1	Ensure separate partition exists for /home	☐	☐
1.1.2.3.2	Ensure nodev option set on /home partition	☐	☐
1.1.2.3.3	Ensure nosuid option set on /home partition	☐	☐
1.1.2.4.1	Ensure separate partition exists for /var	☐	☐
1.1.2.4.2	Ensure nodev option set on /var partition	☐	☐
1.1.2.4.3	Ensure nosuid option set on /var partition	☐	☐
1.1.2.5.2	Ensure nodev option set on /var/tmp partition	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
1.1.2.5.3	Ensure nosuid option set on /var/tmp partition	☐	☐
1.1.2.5.4	Ensure noexec option set on /var/tmp partition	☐	☐
1.1.2.6.1	Ensure separate partition exists for /var/log	☐	☐
1.1.2.6.2	Ensure nodev option set on /var/log partition	☐	☐
1.1.2.6.3	Ensure nosuid option set on /var/log partition	☐	☐
1.1.2.6.4	Ensure noexec option set on /var/log partition	☐	☐
1.1.2.7.1	Ensure separate partition exists for /var/log/audit	☐	☐
1.1.2.7.2	Ensure nodev option set on /var/log/audit partition	☐	☐
1.1.2.7.3	Ensure nosuid option set on /var/log/audit partition	☐	☐
1.1.2.7.4	Ensure noexec option set on /var/log/audit partition	☐	☐
1.2.1.1	Ensure the source.list and .source files use the Signed-By option	☐	☐
1.2.1.2	Ensure weak dependencies are configured	☐	☐
1.2.1.3	Ensure access to gpg key files are configured	☐	☐
1.2.1.4	Ensure access to /etc/apt/trusted.gpg.d directory is configured	☐	☐
1.2.1.5	Ensure access to /etc/apt/auth.conf.d directory is configured	☐	☐
1.2.1.6	Ensure access to files in the /etc/apt/auth.conf.d/ directory is configured	☐	☐
1.2.1.7	Ensure access to /usr/share/keyrings directory is configured	☐	☐
1.2.1.8	Ensure access to /etc/apt/sources.list.d directory is configured	☐	☐
1.2.1.9	Ensure access to files in /etc/apt/sources.list.d are configured	☐	☐
1.2.1.10	Ensure repositories are set to https in /etc/apt/sources.list	☐	☐
1.2.1.11	Ensure repositories are set to https in /etc/apt/sources.list.d directory	☐	☐
1.2.1.12	Ensure AllowInsecureRepositories is configured	☐	☐
1.2.1.13	Ensure AllowWeakRepositories is configured	☐	☐
1.2.1.14	Ensure AllowDowngradeToInsecureRepositories is configured	☐	☐
1.2.1.15	Ensure Acquire::Check-Date is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
1.2.2.1	Ensure updates, patches, and additional security software are installed	☐	☐
1.3.1.1	Ensure apparmor packages are installed	☐	☐
1.3.1.2	Ensure AppArmor is enabled	☐	☐
1.3.1.3	Ensure all AppArmor Profiles are enforcing	☐	☐
1.3.1.4	Ensure kernel.apparmor_restrict_unprivileged_unconfined is enabled	☐	☐
1.4.1	Ensure bootloader password is set	☐	☐
1.4.2	Ensure access to bootloader config is configured	☐	☐
1.5.3	Ensure kernel.yama.ptrace_scope is configured	☐	☐
1.5.5	Ensure kernel.dmesg_restrict is configured	☐	☐
1.5.7	Ensure Automatic Error Reporting is configured	☐	☐
1.5.8	Ensure kernel.kptr_restrict is configured	☐	☐
1.5.9	Ensure kernel.randomize_va_space is configured	☐	☐
1.6.6	Ensure access to /etc/motd is configured	☐	☐
1.6.7	Ensure access to /etc/issue is configured	☐	☐
1.6.8	Ensure access to /etc/issue.net is configured	☐	☐
1.6.9	Ensure access to pam_motd file is configured	☐	☐
1.6.10	Ensure access to sshd warning banner is configured	☐	☐
1.6.11	Ensure update-notifier-motd.service is not in use	☐	☐
1.6.12	Ensure update-notifier-motd.timer is not in use	☐	☐
1.7.2	Ensure GDM disable-user-list is configured	☐	☐
1.7.3	Ensure GDM screen lock is configured	☐	☐
1.7.4	Ensure GDM automount is configured	☐	☐
1.7.5	Ensure GDM autorun-never is configured	☐	☐
1.7.6	Ensure XDMCP is not enabled	☐	☐
1.7.7	Ensure Xwayland is configured	☐	☐
2.1.1	Ensure autofs services are not in use	☐	☐
2.1.2	Ensure avahi daemon services are not in use	☐	☐
2.1.3	Ensure dhcp server services are not in use	☐	☐
2.1.4	Ensure dns server services are not in use	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
2.1.5	Ensure dnsmasq services are not in use	☐	☐
2.1.6	Ensure ftp server services are not in use	☐	☐
2.1.7	Ensure ldap server services are not in use	☐	☐
2.1.8	Ensure message access server services are not in use	☐	☐
2.1.9	Ensure network file system services are not in use	☐	☐
2.1.10	Ensure nis server services are not in use	☐	☐
2.1.11	Ensure print server services are not in use	☐	☐
2.1.12	Ensure rpcbind services are not in use	☐	☐
2.1.13	Ensure rsync services are not in use	☐	☐
2.1.14	Ensure samba file server services are not in use	☐	☐
2.1.15	Ensure snmp services are not in use	☐	☐
2.1.16	Ensure telnet server services are not in use	☐	☐
2.1.17	Ensure tftp server services are not in use	☐	☐
2.1.18	Ensure web proxy server services are not in use	☐	☐
2.1.19	Ensure web server services are not in use	☐	☐
2.1.20	Ensure xinetd services are not in use	☐	☐
2.1.21	Ensure X window server services are not in use	☐	☐
2.1.22	Ensure mail transfer agents are configured for local-only mode	☐	☐
2.1.23	Ensure only approved services are listening on a network interface	☐	☐
2.2.1	Ensure nis client is not installed	☐	☐
2.2.2	Ensure rsh client is not installed	☐	☐
2.2.3	Ensure talk client is not installed	☐	☐
2.2.4	Ensure telnet client is not installed	☐	☐
2.2.5	Ensure ldap client is not installed	☐	☐
2.2.6	Ensure ftp client is not installed	☐	☐
2.3.1.1	Ensure a single time synchronization daemon is in use	☐	☐
2.3.2.1	Ensure systemd-timesyncd configured with authorized timeserver	☐	☐
2.3.2.2	Ensure systemd-timesyncd is enabled and running	☐	☐
2.3.3.1	Ensure chrony is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
2.3.3.3	Ensure chrony is enabled and running	☐	☐
2.4.1.2	Ensure access to /etc/crontab is configured	☐	☐
2.4.1.3	Ensure access to /etc/cron.hourly is configured	☐	☐
2.4.1.4	Ensure access to /etc/cron.daily is configured	☐	☐
2.4.1.5	Ensure access to /etc/cron.weekly is configured	☐	☐
2.4.1.6	Ensure access to /etc/cron.monthly is configured	☐	☐
2.4.1.7	Ensure access to /etc/cron.yearly is configured	☐	☐
2.4.1.8	Ensure access to /etc/cron.d is configured	☐	☐
2.4.1.9	Ensure access to crontab is configured	☐	☐
2.4.2.1	Ensure access to at is configured	☐	☐
3.1.1	Ensure IPv6 status is identified	☐	☐
3.1.3	Ensure bluetooth services are not in use	☐	☐
3.2.1	Ensure atm kernel module is not available	☐	☐
3.2.2	Ensure can kernel module is not available	☐	☐
3.2.3	Ensure dccp kernel module is not available	☐	☐
3.2.4	Ensure rds kernel module is not available	☐	☐
3.2.5	Ensure sctp kernel module is not available	☐	☐
3.2.6	Ensure tipc kernel module is not available	☐	☐
3.2.7	Ensure unneeded network protocol kernel modules are not available	☐	☐
3.3.1.1	Ensure net.ipv4.ip_forward is configured	☐	☐
3.3.1.2	Ensure net.ipv4.conf.all.forwarding is configured	☐	☐
3.3.1.3	Ensure net.ipv4.conf.default.forwarding is configured	☐	☐
3.3.1.4	Ensure net.ipv4.conf.all.send_redirects is configured	☐	☐
3.3.1.5	Ensure net.ipv4.conf.default.send_redirects is configured	☐	☐
3.3.1.6	Ensure net.ipv4.icmp_ignore_bogus_error_responses is configured	☐	☐
3.3.1.7	Ensure net.ipv4.icmp_echo_ignore_broadcasts is configured	☐	☐
3.3.1.8	Ensure net.ipv4.conf.all.accept_redirects is configured	☐	☐
3.3.1.9	Ensure net.ipv4.conf.default.accept_redirects is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
3.3.1.10	Ensure net.ipv4.conf.all.secure_redirects is configured	☐	☐
3.3.1.11	Ensure net.ipv4.conf.default.secure_redirects is configured	☐	☐
3.3.1.12	Ensure net.ipv4.conf.all.rp_filter is configured	☐	☐
3.3.1.13	Ensure net.ipv4.conf.default.rp_filter is configured	☐	☐
3.3.1.14	Ensure net.ipv4.conf.all.accept_source_route is configured	☐	☐
3.3.1.15	Ensure net.ipv4.conf.default.accept_source_route is configured	☐	☐
3.3.1.16	Ensure net.ipv4.conf.all.log_martians is configured	☐	☐
3.3.1.17	Ensure net.ipv4.conf.default.log_martians is configured	☐	☐
3.3.1.18	Ensure net.ipv4.tcp_syncookies is configured	☐	☐
3.3.1.19	Ensure net.ipv4.conf.all.route_localnet is configured	☐	☐
3.3.2.1	Ensure net.ipv6.conf.all.forwarding is configured	☐	☐
3.3.2.2	Ensure net.ipv6.conf.default.forwarding is configured	☐	☐
3.3.2.3	Ensure net.ipv6.conf.all.accept_redirects is configured	☐	☐
3.3.2.4	Ensure net.ipv6.conf.default.accept_redirects is configured	☐	☐
3.3.2.5	Ensure net.ipv6.conf.all.accept_source_route is configured	☐	☐
3.3.2.6	Ensure net.ipv6.conf.default.accept_source_route is configured	☐	☐
3.3.2.7	Ensure net.ipv6.conf.all.accept_ra is configured	☐	☐
3.3.2.8	Ensure net.ipv6.conf.default.accept_ra is configured	☐	☐
4.1.1	Ensure ufw is installed	☐	☐
4.1.2	Ensure ufw service is configured	☐	☐
4.1.3	Ensure ufw incoming default is configured	☐	☐
4.1.4	Ensure ufw outgoing default is configured	☐	☐
4.1.5	Ensure ufw routed default is configured	☐	☐
5.1.1	Ensure access to /etc/ssh/sshd_config is configured	☐	☐
5.1.2	Ensure access to /etc/ssh/sshd_config.d/ directory is configured	☐	☐
5.1.3	Ensure access to SSH private host key files is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
5.1.4	Ensure access to SSH public host key files is configured	☐	☐
5.1.5	Ensure sshd access is configured	☐	☐
5.1.7	Ensure sshd Ciphers are configured	☐	☐
5.1.9	Ensure sshd DisableForwarding is enabled	☐	☐
5.1.10	Ensure sshd GSSAPIAuthentication is disabled	☐	☐
5.1.11	Ensure sshd HostbasedAuthentication is disabled	☐	☐
5.1.12	Ensure sshd IgnoreRhosts is enabled	☐	☐
5.1.13	Ensure sshd KexAlgorithms is configured	☐	☐
5.1.14	Ensure sshd post-quantum cryptography key exchange algorithms are configured	☐	☐
5.1.16	Ensure sshd LogLevel is configured	☐	☐
5.1.17	Ensure sshd MACs are configured	☐	☐
5.1.21	Ensure sshd PermitEmptyPasswords is disabled	☐	☐
5.1.22	Ensure sshd PermitRootLogin is disabled	☐	☐
5.1.24	Ensure sshd UsePAM is enabled	☐	☐
5.2.1	Ensure sudo is installed	☐	☐
5.2.2	Ensure sudo commands use pty	☐	☐
5.2.3	Ensure sudo log file exists	☐	☐
5.2.4	Ensure users must provide password for escalation	☐	☐
5.2.5	Ensure re-authentication for privilege escalation is not disabled globally	☐	☐
5.2.6	Ensure sudo timestamp_timeout is configured	☐	☐
5.2.7	Ensure access to the su command is restricted	☐	☐
5.3.1.3	Ensure latest version of libpam-pwquality is installed	☐	☐
5.3.1.4	Ensure latest version of cracklib-runtime is installed	☐	☐
5.3.2.1	Ensure pam_unix module is enabled	☐	☐
5.3.2.2	Ensure pam_faillock module is enabled	☐	☐
5.3.2.3	Ensure pam_pwquality module is enabled	☐	☐
5.3.2.4	Ensure pam_pwhistory module is enabled	☐	☐
5.3.3.1.1	Ensure password failed attempts lockout is configured	☐	☐
5.3.3.1.2	Ensure password unlock time is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
5.3.3.1.3	Ensure password failed attempts lockout includes root account	☐	☐
5.3.3.2.1	Ensure password number of changed characters is configured	☐	☐
5.3.3.2.2	Ensure password length is configured	☐	☐
5.3.3.2.3	Ensure password complexity is configured	☐	☐
5.3.3.2.4	Ensure password same consecutive characters is configured	☐	☐
5.3.3.2.5	Ensure password maximum sequential characters is configured	☐	☐
5.3.3.2.6	Ensure password dictionary check is enabled	☐	☐
5.3.3.2.7	Ensure password quality checking is enforced	☐	☐
5.3.3.2.8	Ensure password quality is enforced for the root user	☐	☐
5.3.3.3.1	Ensure password history remember is configured	☐	☐
5.3.3.3.2	Ensure password history is enforced for the root user	☐	☐
5.3.3.3.3	Ensure pam_pwhistory includes use_authtok	☐	☐
5.3.3.4.1	Ensure pam_unix does not include nullok	☐	☐
5.3.3.4.2	Ensure pam_unix does not include remember	☐	☐
5.3.3.4.3	Ensure pam_unix includes a strong password hashing algorithm	☐	☐
5.3.3.4.4	Ensure pam_unix includes use_authtok	☐	☐
5.4.1.1	Ensure password expiration is configured	☐	☐
5.4.1.2	Ensure minimum password days is configured	☐	☐
5.4.1.3	Ensure password expiration warning days is configured	☐	☐
5.4.1.4	Ensure strong password hashing algorithm is configured	☐	☐
5.4.1.5	Ensure inactive password lock is configured	☐	☐
5.4.1.6	Ensure all users last password change date is in the past	☐	☐
5.4.2.1	Ensure root is the only UID 0 account	☐	☐
5.4.2.2	Ensure root is the only GID 0 account	☐	☐
5.4.2.3	Ensure group root is the only GID 0 group	☐	☐
5.4.2.4	Ensure root account access is controlled	☐	☐
5.4.2.6	Ensure root user umask is configured	☐	☐
5.4.2.7	Ensure system accounts do not have a valid login shell	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
5.4.2.8	Ensure accounts without a valid login shell are locked	☐	☐
5.4.3.2	Ensure default user shell timeout is configured	☐	☐
5.4.3.3	Ensure default user umask is configured	☐	☐
6.1.1.1.1	Ensure journald service is active	☐	☐
6.1.1.1.2	Ensure systemd-journal-remote service is not in use	☐	☐
6.1.1.1.3	Ensure journald is configured to send logs to rsyslog	☐	☐
6.1.1.1.4	Ensure journald log file access is configured	☐	☐
6.1.1.1.5	Ensure journald log file rotation is configured	☐	☐
6.1.1.1.6	Ensure journald Storage is configured	☐	☐
6.1.1.1.7	Ensure journald Compress is configured	☐	☐
6.1.2.1	Ensure rsyslog is installed	☐	☐
6.1.2.2	Ensure rsyslog service is enabled and active	☐	☐
6.1.2.3	Ensure rsyslog log file creation mode is configured	☐	☐
6.1.2.4	Ensure rsyslog logging is configured	☐	☐
6.1.2.5	Ensure rsyslog is configured to send logs to a remote log host	☐	☐
6.1.2.6	Ensure rsyslog is not configured to receive logs from a remote client	☐	☐
6.1.2.7	Ensure logrotate is configured	☐	☐
6.1.2.8	Ensure rsyslog-gnutls is installed	☐	☐
6.1.2.9	Ensure rsyslog forwarding uses gtls	☐	☐
6.1.2.10	Ensure rsyslog CA certificates are configured	☐	☐
6.1.3.1	Ensure access to all logfiles has been configured	☐	☐
6.2.1.1	Ensure auditd packages are installed	☐	☐
6.2.1.2	Ensure auditd service is enabled and active	☐	☐
6.2.1.3	Ensure auditing for processes that start prior to auditd is enabled	☐	☐
6.2.1.4	Ensure audit_backlog_limit is configured	☐	☐
6.2.2.1	Ensure audit log storage size is configured	☐	☐
6.2.2.2	Ensure audit logs are not automatically deleted	☐	☐
6.2.2.3	Ensure system is disabled when audit logs are full	☐	☐
6.2.2.4	Ensure system warns when audit logs are low on space	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
6.2.3.1	Ensure modification of the /etc/sudoers file is collected	☐	☐
6.2.3.2	Ensure actions as another user are always logged	☐	☐
6.2.3.3	Ensure events that modify the sudo log file are collected	☐	☐
6.2.3.4	Ensure events that modify date and time information are collected	☐	☐
6.2.3.5	Ensure events that modify sethostname and setdomainname are collected	☐	☐
6.2.3.6	Ensure events that modify /etc/issue and /etc/issue.net are collected	☐	☐
6.2.3.7	Ensure events that modify /etc/hosts and /etc/hostname are collected	☐	☐
6.2.3.8	Ensure events that modify the system's network environment are collected	☐	☐
6.2.3.9	Ensure events that modify /etc/NetworkManager directory are collected	☐	☐
6.2.3.10	Ensure use of privileged commands are collected	☐	☐
6.2.3.12	Ensure events that modify /etc/group information are collected	☐	☐
6.2.3.13	Ensure events that modify /etc/passwd information are collected	☐	☐
6.2.3.14	Ensure events that modify /etc/shadow and /etc/gshadow are collected	☐	☐
6.2.3.15	Ensure events that modify /etc/security/opasswd are collected	☐	☐
6.2.3.16	Ensure events that modify /etc/nsswitch.conf file are collected	☐	☐
6.2.3.17	Ensure events that modify /etc/pam.conf and /etc/pam.d/ information are collected	☐	☐
6.2.3.18	Ensure discretionary access control permission modification events chmod,fchmod,fchmodat,fchmodat2 are collected	☐	☐
6.2.3.19	Ensure discretionary access control permission modification events chown,fchown,lchown,fchownat are collected	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
6.2.3.20	Ensure discretionary access control permission modification events setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr collected	☐	☐
6.2.3.21	Ensure successful file system mounts are collected	☐	☐
6.2.3.22	Ensure session initiation information is collected	☐	☐
6.2.3.23	Ensure login and logout events are collected	☐	☐
6.2.3.24	Ensure unlink file deletion events by users are collected	☐	☐
6.2.3.25	Ensure rename file deletion events by users are collected	☐	☐
6.2.3.26	Ensure events that modify the system's Mandatory Access Controls are collected	☐	☐
6.2.3.27	Ensure successful and unsuccessful attempts to use the chcon command are collected	☐	☐
6.2.3.28	Ensure successful and unsuccessful attempts to use the setfacl command are collected	☐	☐
6.2.3.29	Ensure successful and unsuccessful attempts to use the chacl command are collected	☐	☐
6.2.3.30	Ensure successful and unsuccessful attempts to use the usermod command are collected	☐	☐
6.2.3.31	Ensure kernel module loading unloading and modification is collected	☐	☐
6.2.3.32	Ensure kernel "init_module" and "finit_module" loading unloading and modification is collected	☐	☐
6.2.3.33	Ensure kernel "delete_module" loading unloading and modification is collected	☐	☐
6.2.3.34	Ensure the audit configuration is loaded regardless of errors	☐	☐
6.2.3.35	Ensure the audit configuration is immutable	☐	☐
6.2.3.36	Ensure the running and on disk configuration is the same	☐	☐
6.2.4.1	Ensure audit log files mode is configured	☐	☐
6.2.4.2	Ensure audit log files owner is configured	☐	☐
6.2.4.3	Ensure audit log files group owner is configured	☐	☐
6.2.4.4	Ensure the audit log file directory mode is configured	☐	☐
6.2.4.5	Ensure audit configuration files mode is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
6.2.4.6	Ensure audit configuration files owner is configured	☐	☐
6.2.4.7	Ensure audit configuration files group owner is configured	☐	☐
6.2.4.8	Ensure audit tools mode is configured	☐	☐
6.2.4.9	Ensure audit tools owner is configured	☐	☐
6.2.4.10	Ensure audit tools group owner is configured	☐	☐
7.1.1	Ensure access to /etc/passwd is configured	☐	☐
7.1.2	Ensure access to /etc/passwd- is configured	☐	☐
7.1.3	Ensure access to /etc/group is configured	☐	☐
7.1.4	Ensure access to /etc/group- is configured	☐	☐
7.1.5	Ensure access to /etc/shadow is configured	☐	☐
7.1.6	Ensure access to /etc/shadow- is configured	☐	☐
7.1.7	Ensure access to /etc/gshadow is configured	☐	☐
7.1.8	Ensure access to /etc/gshadow- is configured	☐	☐
7.1.9	Ensure access to /etc/shells is configured	☐	☐
7.1.10	Ensure access to /etc/security/opasswd is configured	☐	☐
7.1.11	Ensure world writable files and directories are secured	☐	☐
7.1.12	Ensure no files or directories without an owner and a group exist	☐	☐
7.1.13	Ensure SUID and SGID files are reviewed	☐	☐
7.2.1	Ensure accounts in /etc/passwd use shadowed passwords	☐	☐
7.2.2	Ensure /etc/shadow password fields are not empty	☐	☐
7.2.3	Ensure all groups in /etc/passwd exist in /etc/group	☐	☐
7.2.4	Ensure shadow group is empty	☐	☐
7.2.9	Ensure local interactive user home directories are configured	☐	☐
7.2.10	Ensure local interactive user dot files access is configured	☐	☐
7.2.11	Ensure local interactive user dot directories access is configured	☐	☐

Internal Only - General

Internal Only - General

Appendix: CIS Controls v7 IG 3 Mapped Recommendations

Recommendation		Set Correctly	
		Yes	No
1.1.1.1	Ensure cramfs kernel module is not available	☐	☐
1.1.1.2	Ensure freevxfs kernel module is not available	☐	☐
1.1.1.3	Ensure hfs kernel module is not available	☐	☐
1.1.1.4	Ensure hfsplus kernel module is not available	☐	☐
1.1.1.5	Ensure jffs2 kernel module is not available	☐	☐
1.1.1.6	Ensure overlay kernel module is not available	☐	☐
1.1.1.7	Ensure squashfs kernel module is not available	☐	☐
1.1.1.8	Ensure udf kernel module is not available	☐	☐
1.1.1.9	Ensure firewire-core kernel module is not available	☐	☐
1.1.1.10	Ensure usb-storage kernel module is not available	☐	☐
1.1.1.11	Ensure unused filesystems kernel modules are not available	☐	☐
1.1.2.1.1	Ensure /tmp is tmpfs or a separate partition	☐	☐
1.1.2.1.2	Ensure nodev option set on /tmp partition	☐	☐
1.1.2.1.3	Ensure nosuid option set on /tmp partition	☐	☐
1.1.2.1.4	Ensure noexec option set on /tmp partition	☐	☐
1.1.2.2.1	Ensure /dev/shm is tmpfs or a separate partition	☐	☐
1.1.2.2.2	Ensure nodev option set on /dev/shm partition	☐	☐
1.1.2.2.3	Ensure nosuid option set on /dev/shm partition	☐	☐
1.1.2.2.4	Ensure noexec option set on /dev/shm partition	☐	☐
1.1.2.3.1	Ensure separate partition exists for /home	☐	☐
1.1.2.3.2	Ensure nodev option set on /home partition	☐	☐
1.1.2.3.3	Ensure nosuid option set on /home partition	☐	☐
1.1.2.4.1	Ensure separate partition exists for /var	☐	☐
1.1.2.4.2	Ensure nodev option set on /var partition	☐	☐
1.1.2.4.3	Ensure nosuid option set on /var partition	☐	☐
1.1.2.5.2	Ensure nodev option set on /var/tmp partition	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
1.1.2.5.3	Ensure nosuid option set on /var/tmp partition	☐	☐
1.1.2.5.4	Ensure noexec option set on /var/tmp partition	☐	☐
1.1.2.6.1	Ensure separate partition exists for /var/log	☐	☐
1.1.2.6.2	Ensure nodev option set on /var/log partition	☐	☐
1.1.2.6.3	Ensure nosuid option set on /var/log partition	☐	☐
1.1.2.6.4	Ensure noexec option set on /var/log partition	☐	☐
1.1.2.7.1	Ensure separate partition exists for /var/log/audit	☐	☐
1.1.2.7.2	Ensure nodev option set on /var/log/audit partition	☐	☐
1.1.2.7.3	Ensure nosuid option set on /var/log/audit partition	☐	☐
1.1.2.7.4	Ensure noexec option set on /var/log/audit partition	☐	☐
1.2.1.1	Ensure the source.list and .source files use the Signed-By option	☐	☐
1.2.1.2	Ensure weak dependencies are configured	☐	☐
1.2.1.3	Ensure access to gpg key files are configured	☐	☐
1.2.1.4	Ensure access to /etc/apt/trusted.gpg.d directory is configured	☐	☐
1.2.1.5	Ensure access to /etc/apt/auth.conf.d directory is configured	☐	☐
1.2.1.6	Ensure access to files in the /etc/apt/auth.conf.d/ directory is configured	☐	☐
1.2.1.7	Ensure access to /usr/share/keyrings directory is configured	☐	☐
1.2.1.8	Ensure access to /etc/apt/sources.list.d directory is configured	☐	☐
1.2.1.9	Ensure access to files in /etc/apt/sources.list.d are configured	☐	☐
1.2.1.10	Ensure repositories are set to https in /etc/apt/sources.list	☐	☐
1.2.1.11	Ensure repositories are set to https in /etc/apt/sources.list.d directory	☐	☐
1.2.1.12	Ensure AllowInsecureRepositories is configured	☐	☐
1.2.1.13	Ensure AllowWeakRepositories is configured	☐	☐
1.2.1.14	Ensure AllowDowngradeToInsecureRepositories is configured	☐	☐
1.2.1.15	Ensure Acquire::Check-Date is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
1.2.2.1	Ensure updates, patches, and additional security software are installed	☐	☐
1.3.1.1	Ensure apparmor packages are installed	☐	☐
1.3.1.2	Ensure AppArmor is enabled	☐	☐
1.3.1.3	Ensure all AppArmor Profiles are enforcing	☐	☐
1.3.1.4	Ensure kernel.apparmor_restrict_unprivileged_unconfined is enabled	☐	☐
1.4.1	Ensure bootloader password is set	☐	☐
1.4.2	Ensure access to bootloader config is configured	☐	☐
1.5.3	Ensure kernel.yama.ptrace_scope is configured	☐	☐
1.5.5	Ensure kernel.dmesg_restrict is configured	☐	☐
1.5.6	Ensure prelink is not installed	☐	☐
1.5.7	Ensure Automatic Error Reporting is configured	☐	☐
1.5.8	Ensure kernel.kptr_restrict is configured	☐	☐
1.5.9	Ensure kernel.randomize_va_space is configured	☐	☐
1.6.6	Ensure access to /etc/motd is configured	☐	☐
1.6.7	Ensure access to /etc/issue is configured	☐	☐
1.6.8	Ensure access to /etc/issue.net is configured	☐	☐
1.6.9	Ensure access to pam_motd file is configured	☐	☐
1.6.10	Ensure access to sshd warning banner is configured	☐	☐
1.6.11	Ensure update-notifier-motd.service is not in use	☐	☐
1.6.12	Ensure update-notifier-motd.timer is not in use	☐	☐
1.7.2	Ensure GDM disable-user-list is configured	☐	☐
1.7.3	Ensure GDM screen lock is configured	☐	☐
1.7.4	Ensure GDM automount is configured	☐	☐
1.7.5	Ensure GDM autorun-never is configured	☐	☐
1.7.6	Ensure XDMCP is not enabled	☐	☐
1.7.7	Ensure Xwayland is configured	☐	☐
2.1.1	Ensure autofs services are not in use	☐	☐
2.1.2	Ensure avahi daemon services are not in use	☐	☐
2.1.3	Ensure dhcp server services are not in use	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
2.1.4	Ensure dns server services are not in use	☐	☐
2.1.5	Ensure dnsmasq services are not in use	☐	☐
2.1.6	Ensure ftp server services are not in use	☐	☐
2.1.7	Ensure ldap server services are not in use	☐	☐
2.1.8	Ensure message access server services are not in use	☐	☐
2.1.9	Ensure network file system services are not in use	☐	☐
2.1.10	Ensure nis server services are not in use	☐	☐
2.1.11	Ensure print server services are not in use	☐	☐
2.1.12	Ensure rpcbind services are not in use	☐	☐
2.1.13	Ensure rsync services are not in use	☐	☐
2.1.14	Ensure samba file server services are not in use	☐	☐
2.1.15	Ensure snmp services are not in use	☐	☐
2.1.16	Ensure telnet server services are not in use	☐	☐
2.1.17	Ensure tftp server services are not in use	☐	☐
2.1.18	Ensure web proxy server services are not in use	☐	☐
2.1.19	Ensure web server services are not in use	☐	☐
2.1.20	Ensure xinetd services are not in use	☐	☐
2.1.21	Ensure X window server services are not in use	☐	☐
2.1.22	Ensure mail transfer agents are configured for local-only mode	☐	☐
2.1.23	Ensure only approved services are listening on a network interface	☐	☐
2.2.1	Ensure nis client is not installed	☐	☐
2.2.2	Ensure rsh client is not installed	☐	☐
2.2.3	Ensure talk client is not installed	☐	☐
2.2.4	Ensure telnet client is not installed	☐	☐
2.2.5	Ensure ldap client is not installed	☐	☐
2.2.6	Ensure ftp client is not installed	☐	☐
2.3.1.1	Ensure a single time synchronization daemon is in use	☐	☐
2.3.2.1	Ensure systemd-timesyncd configured with authorized timeserver	☐	☐
2.3.2.2	Ensure systemd-timesyncd is enabled and running	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
2.3.3.1	Ensure chrony is configured	☐	☐
2.3.3.3	Ensure chrony is enabled and running	☐	☐
2.4.1.2	Ensure access to /etc/crontab is configured	☐	☐
2.4.1.3	Ensure access to /etc/cron.hourly is configured	☐	☐
2.4.1.4	Ensure access to /etc/cron.daily is configured	☐	☐
2.4.1.5	Ensure access to /etc/cron.weekly is configured	☐	☐
2.4.1.6	Ensure access to /etc/cron.monthly is configured	☐	☐
2.4.1.7	Ensure access to /etc/cron.yearly is configured	☐	☐
2.4.1.8	Ensure access to /etc/cron.d is configured	☐	☐
2.4.1.9	Ensure access to crontab is configured	☐	☐
2.4.2.1	Ensure access to at is configured	☐	☐
3.1.1	Ensure IPv6 status is identified	☐	☐
3.1.2	Ensure wireless interfaces are not available	☐	☐
3.1.3	Ensure bluetooth services are not in use	☐	☐
3.2.1	Ensure atm kernel module is not available	☐	☐
3.2.2	Ensure can kernel module is not available	☐	☐
3.2.3	Ensure dccp kernel module is not available	☐	☐
3.2.4	Ensure rds kernel module is not available	☐	☐
3.2.5	Ensure sctp kernel module is not available	☐	☐
3.2.6	Ensure tipc kernel module is not available	☐	☐
3.2.7	Ensure unneeded network protocol kernel modules are not available	☐	☐
3.3.1.1	Ensure net.ipv4.ip_forward is configured	☐	☐
3.3.1.2	Ensure net.ipv4.conf.all.forwarding is configured	☐	☐
3.3.1.3	Ensure net.ipv4.conf.default.forwarding is configured	☐	☐
3.3.1.4	Ensure net.ipv4.conf.all.send_redirects is configured	☐	☐
3.3.1.5	Ensure net.ipv4.conf.default.send_redirects is configured	☐	☐
3.3.1.6	Ensure net.ipv4.icmp_ignore_bogus_error_responses is configured	☐	☐
3.3.1.7	Ensure net.ipv4.icmp_echo_ignore_broadcasts is configured	☐	☐
3.3.1.8	Ensure net.ipv4.conf.all.accept_redirects is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
3.3.1.9	Ensure net.ipv4.conf.default.accept_redirects is configured	☐	☐
3.3.1.10	Ensure net.ipv4.conf.all.secure_redirects is configured	☐	☐
3.3.1.11	Ensure net.ipv4.conf.default.secure_redirects is configured	☐	☐
3.3.1.12	Ensure net.ipv4.conf.all.rp_filter is configured	☐	☐
3.3.1.13	Ensure net.ipv4.conf.default.rp_filter is configured	☐	☐
3.3.1.14	Ensure net.ipv4.conf.all.accept_source_route is configured	☐	☐
3.3.1.15	Ensure net.ipv4.conf.default.accept_source_route is configured	☐	☐
3.3.1.16	Ensure net.ipv4.conf.all.log_martians is configured	☐	☐
3.3.1.17	Ensure net.ipv4.conf.default.log_martians is configured	☐	☐
3.3.1.18	Ensure net.ipv4.tcp_syncookies is configured	☐	☐
3.3.1.19	Ensure net.ipv4.conf.all.route_localnet is configured	☐	☐
3.3.2.1	Ensure net.ipv6.conf.all.forwarding is configured	☐	☐
3.3.2.2	Ensure net.ipv6.conf.default.forwarding is configured	☐	☐
3.3.2.3	Ensure net.ipv6.conf.all.accept_redirects is configured	☐	☐
3.3.2.4	Ensure net.ipv6.conf.default.accept_redirects is configured	☐	☐
3.3.2.5	Ensure net.ipv6.conf.all.accept_source_route is configured	☐	☐
3.3.2.6	Ensure net.ipv6.conf.default.accept_source_route is configured	☐	☐
3.3.2.7	Ensure net.ipv6.conf.all.accept_ra is configured	☐	☐
3.3.2.8	Ensure net.ipv6.conf.default.accept_ra is configured	☐	☐
4.1.1	Ensure ufw is installed	☐	☐
4.1.2	Ensure ufw service is configured	☐	☐
4.1.3	Ensure ufw incoming default is configured	☐	☐
4.1.4	Ensure ufw outgoing default is configured	☐	☐
4.1.5	Ensure ufw routed default is configured	☐	☐
5.1.1	Ensure access to /etc/ssh/sshd_config is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
5.1.2	Ensure access to /etc/ssh/sshd_config.d/ directory is configured	☐	☐
5.1.3	Ensure access to SSH private host key files is configured	☐	☐
5.1.4	Ensure access to SSH public host key files is configured	☐	☐
5.1.5	Ensure sshd access is configured	☐	☐
5.1.7	Ensure sshd Ciphers are configured	☐	☐
5.1.9	Ensure sshd DisableForwarding is enabled	☐	☐
5.1.10	Ensure sshd GSSAPIAuthentication is disabled	☐	☐
5.1.11	Ensure sshd HostbasedAuthentication is disabled	☐	☐
5.1.12	Ensure sshd IgnoreRhosts is enabled	☐	☐
5.1.13	Ensure sshd KexAlgorithms is configured	☐	☐
5.1.14	Ensure sshd post-quantum cryptography key exchange algorithms are configured	☐	☐
5.1.16	Ensure sshd LogLevel is configured	☐	☐
5.1.17	Ensure sshd MACs are configured	☐	☐
5.1.18	Ensure sshd MaxAuthTries is configured	☐	☐
5.1.21	Ensure sshd PermitEmptyPasswords is disabled	☐	☐
5.1.22	Ensure sshd PermitRootLogin is disabled	☐	☐
5.1.24	Ensure sshd UsePAM is enabled	☐	☐
5.2.1	Ensure sudo is installed	☐	☐
5.2.2	Ensure sudo commands use pty	☐	☐
5.2.3	Ensure sudo log file exists	☐	☐
5.2.4	Ensure users must provide password for escalation	☐	☐
5.2.5	Ensure re-authentication for privilege escalation is not disabled globally	☐	☐
5.2.6	Ensure sudo timestamp_timeout is configured	☐	☐
5.2.7	Ensure access to the su command is restricted	☐	☐
5.3.1.3	Ensure latest version of libpam-pwquality is installed	☐	☐
5.3.1.4	Ensure latest version of cracklib-runtime is installed	☐	☐
5.3.2.1	Ensure pam_unix module is enabled	☐	☐
5.3.2.2	Ensure pam_faillock module is enabled	☐	☐
5.3.2.3	Ensure pam_pwquality module is enabled	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
5.3.2.4	Ensure pam_pwhistory module is enabled	☐	☐
5.3.3.1.1	Ensure password failed attempts lockout is configured	☐	☐
5.3.3.1.2	Ensure password unlock time is configured	☐	☐
5.3.3.1.3	Ensure password failed attempts lockout includes root account	☐	☐
5.3.3.2.1	Ensure password number of changed characters is configured	☐	☐
5.3.3.2.2	Ensure password length is configured	☐	☐
5.3.3.2.3	Ensure password complexity is configured	☐	☐
5.3.3.2.4	Ensure password same consecutive characters is configured	☐	☐
5.3.3.2.5	Ensure password maximum sequential characters is configured	☐	☐
5.3.3.2.6	Ensure password dictionary check is enabled	☐	☐
5.3.3.2.7	Ensure password quality checking is enforced	☐	☐
5.3.3.2.8	Ensure password quality is enforced for the root user	☐	☐
5.3.3.3.1	Ensure password history remember is configured	☐	☐
5.3.3.3.2	Ensure password history is enforced for the root user	☐	☐
5.3.3.3.3	Ensure pam_pwhistory includes use_authtok	☐	☐
5.3.3.4.1	Ensure pam_unix does not include nullok	☐	☐
5.3.3.4.2	Ensure pam_unix does not include remember	☐	☐
5.3.3.4.3	Ensure pam_unix includes a strong password hashing algorithm	☐	☐
5.3.3.4.4	Ensure pam_unix includes use_authtok	☐	☐
5.4.1.1	Ensure password expiration is configured	☐	☐
5.4.1.2	Ensure minimum password days is configured	☐	☐
5.4.1.3	Ensure password expiration warning days is configured	☐	☐
5.4.1.4	Ensure strong password hashing algorithm is configured	☐	☐
5.4.1.5	Ensure inactive password lock is configured	☐	☐
5.4.1.6	Ensure all users last password change date is in the past	☐	☐
5.4.2.1	Ensure root is the only UID 0 account	☐	☐
5.4.2.2	Ensure root is the only GID 0 account	☐	☐
5.4.2.3	Ensure group root is the only GID 0 group	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
5.4.2.4	Ensure root account access is controlled	☐	☐
5.4.2.6	Ensure root user umask is configured	☐	☐
5.4.2.7	Ensure system accounts do not have a valid login shell	☐	☐
5.4.2.8	Ensure accounts without a valid login shell are locked	☐	☐
5.4.3.2	Ensure default user shell timeout is configured	☐	☐
5.4.3.3	Ensure default user umask is configured	☐	☐
6.1.1.1.1	Ensure journald service is active	☐	☐
6.1.1.1.2	Ensure systemd-journal-remote service is not in use	☐	☐
6.1.1.1.3	Ensure journald is configured to send logs to rsyslog	☐	☐
6.1.1.1.4	Ensure journald log file access is configured	☐	☐
6.1.1.1.5	Ensure journald log file rotation is configured	☐	☐
6.1.1.1.6	Ensure journald Storage is configured	☐	☐
6.1.1.1.7	Ensure journald Compress is configured	☐	☐
6.1.2.1	Ensure rsyslog is installed	☐	☐
6.1.2.2	Ensure rsyslog service is enabled and active	☐	☐
6.1.2.3	Ensure rsyslog log file creation mode is configured	☐	☐
6.1.2.4	Ensure rsyslog logging is configured	☐	☐
6.1.2.5	Ensure rsyslog is configured to send logs to a remote log host	☐	☐
6.1.2.6	Ensure rsyslog is not configured to receive logs from a remote client	☐	☐
6.1.2.7	Ensure logrotate is configured	☐	☐
6.1.2.8	Ensure rsyslog-gnutls is installed	☐	☐
6.1.2.9	Ensure rsyslog forwarding uses gtls	☐	☐
6.1.2.10	Ensure rsyslog CA certificates are configured	☐	☐
6.1.3.1	Ensure access to all logfiles has been configured	☐	☐
6.2.1.1	Ensure auditd packages are installed	☐	☐
6.2.1.2	Ensure auditd service is enabled and active	☐	☐
6.2.1.3	Ensure auditing for processes that start prior to auditd is enabled	☐	☐
6.2.1.4	Ensure audit_backlog_limit is configured	☐	☐
6.2.2.1	Ensure audit log storage size is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
6.2.2.2	Ensure audit logs are not automatically deleted	☐	☐
6.2.2.3	Ensure system is disabled when audit logs are full	☐	☐
6.2.2.4	Ensure system warns when audit logs are low on space	☐	☐
6.2.3.1	Ensure modification of the /etc/sudoers file is collected	☐	☐
6.2.3.2	Ensure actions as another user are always logged	☐	☐
6.2.3.3	Ensure events that modify the sudo log file are collected	☐	☐
6.2.3.4	Ensure events that modify date and time information are collected	☐	☐
6.2.3.5	Ensure events that modify sethostname and setdomainname are collected	☐	☐
6.2.3.6	Ensure events that modify /etc/issue and /etc/issue.net are collected	☐	☐
6.2.3.7	Ensure events that modify /etc/hosts and /etc/hostname are collected	☐	☐
6.2.3.8	Ensure events that modify the system's network environment are collected	☐	☐
6.2.3.9	Ensure events that modify /etc/NetworkManager directory are collected	☐	☐
6.2.3.10	Ensure use of privileged commands are collected	☐	☐
6.2.3.11	Ensure unsuccessful file access attempts are collected	☐	☐
6.2.3.12	Ensure events that modify /etc/group information are collected	☐	☐
6.2.3.13	Ensure events that modify /etc/passwd information are collected	☐	☐
6.2.3.14	Ensure events that modify /etc/shadow and /etc/gshadow are collected	☐	☐
6.2.3.15	Ensure events that modify /etc/security/opasswd are collected	☐	☐
6.2.3.16	Ensure events that modify /etc/nsswitch.conf file are collected	☐	☐
6.2.3.17	Ensure events that modify /etc/pam.conf and /etc/pam.d/ information are collected	☐	☐
6.2.3.18	Ensure discretionary access control permission modification events chmod,fchmod,fchmodat,fchmodat2 are collected	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
6.2.3.19	Ensure discretionary access control permission modification events chown,fchown,lchown,fchownat are collected	☐	☐
6.2.3.20	Ensure discretionary access control permission modification events setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr collected	☐	☐
6.2.3.21	Ensure successful file system mounts are collected	☐	☐
6.2.3.22	Ensure session initiation information is collected	☐	☐
6.2.3.23	Ensure login and logout events are collected	☐	☐
6.2.3.24	Ensure unlink file deletion events by users are collected	☐	☐
6.2.3.25	Ensure rename file deletion events by users are collected	☐	☐
6.2.3.26	Ensure events that modify the system's Mandatory Access Controls are collected	☐	☐
6.2.3.27	Ensure successful and unsuccessful attempts to use the chcon command are collected	☐	☐
6.2.3.28	Ensure successful and unsuccessful attempts to use the setfacl command are collected	☐	☐
6.2.3.29	Ensure successful and unsuccessful attempts to use the chacl command are collected	☐	☐
6.2.3.30	Ensure successful and unsuccessful attempts to use the usermod command are collected	☐	☐
6.2.3.31	Ensure kernel module loading unloading and modification is collected	☐	☐
6.2.3.32	Ensure kernel "init_module" and "finit_module" loading unloading and modification is collected	☐	☐
6.2.3.33	Ensure kernel "delete_module" loading unloading and modification is collected	☐	☐
6.2.3.34	Ensure the audit configuration is loaded regardless of errors	☐	☐
6.2.3.35	Ensure the audit configuration is immutable	☐	☐
6.2.3.36	Ensure the running and on disk configuration is the same	☐	☐
6.2.4.1	Ensure audit log files mode is configured	☐	☐
6.2.4.2	Ensure audit log files owner is configured	☐	☐
6.2.4.3	Ensure audit log files group owner is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
6.2.4.4	Ensure the audit log file directory mode is configured	☐	☐
6.2.4.5	Ensure audit configuration files mode is configured	☐	☐
6.2.4.6	Ensure audit configuration files owner is configured	☐	☐
6.2.4.7	Ensure audit configuration files group owner is configured	☐	☐
6.2.4.8	Ensure audit tools mode is configured	☐	☐
6.2.4.9	Ensure audit tools owner is configured	☐	☐
6.2.4.10	Ensure audit tools group owner is configured	☐	☐
6.3.1	Ensure AIDE is installed	☐	☐
6.3.2	Ensure filesystem integrity is regularly checked	☐	☐
7.1.1	Ensure access to /etc/passwd is configured	☐	☐
7.1.2	Ensure access to /etc/passwd- is configured	☐	☐
7.1.3	Ensure access to /etc/group is configured	☐	☐
7.1.4	Ensure access to /etc/group- is configured	☐	☐
7.1.5	Ensure access to /etc/shadow is configured	☐	☐
7.1.6	Ensure access to /etc/shadow- is configured	☐	☐
7.1.7	Ensure access to /etc/gshadow is configured	☐	☐
7.1.8	Ensure access to /etc/gshadow- is configured	☐	☐
7.1.9	Ensure access to /etc/shells is configured	☐	☐
7.1.10	Ensure access to /etc/security/opasswd is configured	☐	☐
7.1.11	Ensure world writable files and directories are secured	☐	☐
7.1.12	Ensure no files or directories without an owner and a group exist	☐	☐
7.1.13	Ensure SUID and SGID files are reviewed	☐	☐
7.2.1	Ensure accounts in /etc/passwd use shadowed passwords	☐	☐
7.2.2	Ensure /etc/shadow password fields are not empty	☐	☐
7.2.3	Ensure all groups in /etc/passwd exist in /etc/group	☐	☐
7.2.4	Ensure shadow group is empty	☐	☐
7.2.9	Ensure local interactive user home directories are configured	☐	☐
7.2.10	Ensure local interactive user dot files access is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
7.2.11	Ensure local interactive user dot directories access is configured	☐	☐

Internal Only - General

Appendix: CIS Controls v7 Unmapped Recommendations

Recommendation		Set Correctly	
		Yes	No
	No unmapped recommendations to CIS Controls v7	☐	☐

Internal Only - General

Appendix: CIS Controls v8 IG 1 Mapped Recommendations

Recommendation		Set Correctly	
		Yes	No
1.1.1.10	Ensure usb-storage kernel module is not available	☐	☐
1.1.2.1.1	Ensure /tmp is tmpfs or a separate partition	☐	☐
1.1.2.1.2	Ensure nodev option set on /tmp partition	☐	☐
1.1.2.1.3	Ensure nosuid option set on /tmp partition	☐	☐
1.1.2.1.4	Ensure noexec option set on /tmp partition	☐	☐
1.1.2.2.1	Ensure /dev/shm is tmpfs or a separate partition	☐	☐
1.1.2.2.2	Ensure nodev option set on /dev/shm partition	☐	☐
1.1.2.2.3	Ensure nosuid option set on /dev/shm partition	☐	☐
1.1.2.2.4	Ensure noexec option set on /dev/shm partition	☐	☐
1.1.2.3.1	Ensure separate partition exists for /home	☐	☐
1.1.2.3.2	Ensure nodev option set on /home partition	☐	☐
1.1.2.3.3	Ensure nosuid option set on /home partition	☐	☐
1.1.2.4.1	Ensure separate partition exists for /var	☐	☐
1.1.2.4.2	Ensure nodev option set on /var partition	☐	☐
1.1.2.4.3	Ensure nosuid option set on /var partition	☐	☐
1.1.2.5.2	Ensure nodev option set on /var/tmp partition	☐	☐
1.1.2.5.3	Ensure nosuid option set on /var/tmp partition	☐	☐
1.1.2.5.4	Ensure noexec option set on /var/tmp partition	☐	☐
1.1.2.6.1	Ensure separate partition exists for /var/log	☐	☐
1.1.2.6.2	Ensure nodev option set on /var/log partition	☐	☐
1.1.2.6.3	Ensure nosuid option set on /var/log partition	☐	☐
1.1.2.6.4	Ensure noexec option set on /var/log partition	☐	☐
1.1.2.7.1	Ensure separate partition exists for /var/log/audit	☐	☐
1.1.2.7.2	Ensure nodev option set on /var/log/audit partition	☐	☐
1.1.2.7.3	Ensure nosuid option set on /var/log/audit partition	☐	☐
1.1.2.7.4	Ensure noexec option set on /var/log/audit partition	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
1.2.1.1	Ensure the source.list and .source files use the Signed-By option	☐	☐
1.2.1.3	Ensure access to gpg key files are configured	☐	☐
1.2.1.4	Ensure access to /etc/apt/trusted.gpg.d directory is configured	☐	☐
1.2.1.5	Ensure access to /etc/apt/auth.conf.d directory is configured	☐	☐
1.2.1.6	Ensure access to files in the /etc/apt/auth.conf.d/ directory is configured	☐	☐
1.2.1.7	Ensure access to /usr/share/keyrings directory is configured	☐	☐
1.2.1.8	Ensure access to /etc/apt/sources.list.d directory is configured	☐	☐
1.2.1.9	Ensure access to files in /etc/apt/sources.list.d are configured	☐	☐
1.2.1.12	Ensure AllowInsecureRepositories is configured	☐	☐
1.2.1.13	Ensure AllowWeakRepositories is configured	☐	☐
1.2.1.14	Ensure AllowDowngradeToInsecureRepositories is configured	☐	☐
1.2.1.15	Ensure Acquire::Check-Date is configured	☐	☐
1.2.2.1	Ensure updates, patches, and additional security software are installed	☐	☐
1.3.1.1	Ensure apparmor packages are installed	☐	☐
1.3.1.2	Ensure AppArmor is enabled	☐	☐
1.3.1.3	Ensure all AppArmor Profiles are enforcing	☐	☐
1.4.1	Ensure bootloader password is set	☐	☐
1.4.2	Ensure access to bootloader config is configured	☐	☐
1.5.5	Ensure kernel.dmesg_restrict is configured	☐	☐
1.5.8	Ensure kernel.kptr_restrict is configured	☐	☐
1.6.6	Ensure access to /etc/motd is configured	☐	☐
1.6.7	Ensure access to /etc/issue is configured	☐	☐
1.6.8	Ensure access to /etc/issue.net is configured	☐	☐
1.6.9	Ensure access to pam_motd file is configured	☐	☐
1.6.10	Ensure access to sshd warning banner is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
1.7.3	Ensure GDM screen lock is configured	☐	☐
1.7.4	Ensure GDM automount is configured	☐	☐
1.7.5	Ensure GDM autorun-never is configured	☐	☐
2.1.1	Ensure autofs services are not in use	☐	☐
2.4.1.2	Ensure access to /etc/crontab is configured	☐	☐
2.4.1.3	Ensure access to /etc/cron.hourly is configured	☐	☐
2.4.1.4	Ensure access to /etc/cron.daily is configured	☐	☐
2.4.1.5	Ensure access to /etc/cron.weekly is configured	☐	☐
2.4.1.6	Ensure access to /etc/cron.monthly is configured	☐	☐
2.4.1.7	Ensure access to /etc/cron.yearly is configured	☐	☐
2.4.1.8	Ensure access to /etc/cron.d is configured	☐	☐
2.4.1.9	Ensure access to crontab is configured	☐	☐
2.4.2.1	Ensure access to at is configured	☐	☐
4.1.1	Ensure ufw is installed	☐	☐
4.1.2	Ensure ufw service is configured	☐	☐
4.1.3	Ensure ufw incoming default is configured	☐	☐
4.1.4	Ensure ufw outgoing default is configured	☐	☐
4.1.5	Ensure ufw routed default is configured	☐	☐
5.1.1	Ensure access to /etc/ssh/sshd_config is configured	☐	☐
5.1.2	Ensure access to /etc/ssh/sshd_config.d/ directory is configured	☐	☐
5.1.3	Ensure access to SSH private host key files is configured	☐	☐
5.1.4	Ensure access to SSH public host key files is configured	☐	☐
5.1.5	Ensure sshd access is configured	☐	☐
5.1.10	Ensure sshd GSSAPIAuthentication is disabled	☐	☐
5.1.11	Ensure sshd HostbasedAuthentication is disabled	☐	☐
5.1.12	Ensure sshd IgnoreRhosts is enabled	☐	☐
5.1.16	Ensure sshd LogLevel is configured	☐	☐
5.1.21	Ensure sshd PermitEmptyPasswords is disabled	☐	☐
5.1.22	Ensure sshd PermitRootLogin is disabled	☐	☐
5.1.24	Ensure sshd UsePAM is enabled	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
5.2.1	Ensure sudo is installed	☐	☐
5.2.2	Ensure sudo commands use pty	☐	☐
5.2.4	Ensure users must provide password for escalation	☐	☐
5.2.5	Ensure re-authentication for privilege escalation is not disabled globally	☐	☐
5.2.6	Ensure sudo timestamp_timeout is configured	☐	☐
5.2.7	Ensure access to the su command is restricted	☐	☐
5.3.1.3	Ensure latest version of libpam-pwquality is installed	☐	☐
5.3.1.4	Ensure latest version of cracklib-runtime is installed	☐	☐
5.3.2.1	Ensure pam_unix module is enabled	☐	☐
5.3.2.2	Ensure pam_faillock module is enabled	☐	☐
5.3.2.3	Ensure pam_pwquality module is enabled	☐	☐
5.3.2.4	Ensure pam_pwhistory module is enabled	☐	☐
5.3.3.1.1	Ensure password failed attempts lockout is configured	☐	☐
5.3.3.1.2	Ensure password unlock time is configured	☐	☐
5.3.3.1.3	Ensure password failed attempts lockout includes root account	☐	☐
5.3.3.2.1	Ensure password number of changed characters is configured	☐	☐
5.3.3.2.2	Ensure password length is configured	☐	☐
5.3.3.2.3	Ensure password complexity is configured	☐	☐
5.3.3.2.4	Ensure password same consecutive characters is configured	☐	☐
5.3.3.2.5	Ensure password maximum sequential characters is configured	☐	☐
5.3.3.2.6	Ensure password dictionary check is enabled	☐	☐
5.3.3.2.7	Ensure password quality checking is enforced	☐	☐
5.3.3.2.8	Ensure password quality is enforced for the root user	☐	☐
5.3.3.3.1	Ensure password history remember is configured	☐	☐
5.3.3.3.2	Ensure password history is enforced for the root user	☐	☐
5.3.3.4.1	Ensure pam_unix does not include nullok	☐	☐
5.3.3.4.2	Ensure pam_unix does not include remember	☐	☐
5.4.1.1	Ensure password expiration is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
5.4.1.2	Ensure minimum password days is configured	☐	☐
5.4.1.3	Ensure password expiration warning days is configured	☐	☐
5.4.1.5	Ensure inactive password lock is configured	☐	☐
5.4.1.6	Ensure all users last password change date is in the past	☐	☐
5.4.2.1	Ensure root is the only UID 0 account	☐	☐
5.4.2.2	Ensure root is the only GID 0 account	☐	☐
5.4.2.3	Ensure group root is the only GID 0 group	☐	☐
5.4.2.4	Ensure root account access is controlled	☐	☐
5.4.2.6	Ensure root user umask is configured	☐	☐
5.4.2.7	Ensure system accounts do not have a valid login shell	☐	☐
5.4.2.8	Ensure accounts without a valid login shell are locked	☐	☐
5.4.3.2	Ensure default user shell timeout is configured	☐	☐
5.4.3.3	Ensure default user umask is configured	☐	☐
6.1.1.1.1	Ensure journald service is active	☐	☐
6.1.1.1.3	Ensure journald is configured to send logs to rsyslog	☐	☐
6.1.1.1.4	Ensure journald log file access is configured	☐	☐
6.1.1.1.5	Ensure journald log file rotation is configured	☐	☐
6.1.1.1.6	Ensure journald Storage is configured	☐	☐
6.1.1.1.7	Ensure journald Compress is configured	☐	☐
6.1.2.1	Ensure rsyslog is installed	☐	☐
6.1.2.2	Ensure rsyslog service is enabled and active	☐	☐
6.1.2.3	Ensure rsyslog log file creation mode is configured	☐	☐
6.1.2.4	Ensure rsyslog logging is configured	☐	☐
6.1.2.5	Ensure rsyslog is configured to send logs to a remote log host	☐	☐
6.1.2.7	Ensure logrotate is configured	☐	☐
6.1.2.8	Ensure rsyslog-gnutls is installed	☐	☐
6.1.3.1	Ensure access to all logfiles has been configured	☐	☐
6.2.1.2	Ensure auditd service is enabled and active	☐	☐
6.2.1.3	Ensure auditing for processes that start prior to auditd is enabled	☐	☐
6.2.1.4	Ensure audit_backlog_limit is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
6.2.2.1	Ensure audit log storage size is configured	☐	☐
6.2.2.2	Ensure audit logs are not automatically deleted	☐	☐
6.2.2.3	Ensure system is disabled when audit logs are full	☐	☐
6.2.2.4	Ensure system warns when audit logs are low on space	☐	☐
6.2.3.27	Ensure successful and unsuccessful attempts to use the chcon command are collected	☐	☐
6.2.3.28	Ensure successful and unsuccessful attempts to use the setfacl command are collected	☐	☐
6.2.3.29	Ensure successful and unsuccessful attempts to use the chacl command are collected	☐	☐
6.2.3.30	Ensure successful and unsuccessful attempts to use the usermod command are collected	☐	☐
6.2.3.35	Ensure the audit configuration is immutable	☐	☐
6.2.4.1	Ensure audit log files mode is configured	☐	☐
6.2.4.2	Ensure audit log files owner is configured	☐	☐
6.2.4.3	Ensure audit log files group owner is configured	☐	☐
6.2.4.4	Ensure the audit log file directory mode is configured	☐	☐
6.2.4.5	Ensure audit configuration files mode is configured	☐	☐
6.2.4.6	Ensure audit configuration files owner is configured	☐	☐
6.2.4.7	Ensure audit configuration files group owner is configured	☐	☐
6.2.4.8	Ensure audit tools mode is configured	☐	☐
6.2.4.9	Ensure audit tools owner is configured	☐	☐
6.2.4.10	Ensure audit tools group owner is configured	☐	☐
7.1.1	Ensure access to /etc/passwd is configured	☐	☐
7.1.2	Ensure access to /etc/passwd- is configured	☐	☐
7.1.3	Ensure access to /etc/group is configured	☐	☐
7.1.4	Ensure access to /etc/group- is configured	☐	☐
7.1.5	Ensure access to /etc/shadow is configured	☐	☐
7.1.6	Ensure access to /etc/shadow- is configured	☐	☐
7.1.7	Ensure access to /etc/gshadow is configured	☐	☐
7.1.8	Ensure access to /etc/gshadow- is configured	☐	☐
7.1.9	Ensure access to /etc/shells is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
7.1.10	Ensure access to /etc/security/opasswd is configured	☐	☐
7.1.11	Ensure world writable files and directories are secured	☐	☐
7.1.12	Ensure no files or directories without an owner and a group exist	☐	☐
7.1.13	Ensure SUID and SGID files are reviewed	☐	☐
7.2.2	Ensure /etc/shadow password fields are not empty	☐	☐
7.2.3	Ensure all groups in /etc/passwd exist in /etc/group	☐	☐
7.2.4	Ensure shadow group is empty	☐	☐
7.2.9	Ensure local interactive user home directories are configured	☐	☐
7.2.10	Ensure local interactive user dot files access is configured	☐	☐
7.2.11	Ensure local interactive user dot directories access is configured	☐	☐

Internal Only - General

Appendix: CIS Controls v8 IG 2 Mapped Recommendations

Recommendation		Set Correctly	
		Yes	No
1.1.1.1	Ensure cramfs kernel module is not available	☐	☐
1.1.1.2	Ensure freevxfs kernel module is not available	☐	☐
1.1.1.3	Ensure hfs kernel module is not available	☐	☐
1.1.1.4	Ensure hfsplus kernel module is not available	☐	☐
1.1.1.5	Ensure jffs2 kernel module is not available	☐	☐
1.1.1.6	Ensure overlay kernel module is not available	☐	☐
1.1.1.7	Ensure squashfs kernel module is not available	☐	☐
1.1.1.8	Ensure udf kernel module is not available	☐	☐
1.1.1.9	Ensure firewire-core kernel module is not available	☐	☐
1.1.1.10	Ensure usb-storage kernel module is not available	☐	☐
1.1.1.11	Ensure unused filesystems kernel modules are not available	☐	☐
1.1.2.1.1	Ensure /tmp is tmpfs or a separate partition	☐	☐
1.1.2.1.2	Ensure nodev option set on /tmp partition	☐	☐
1.1.2.1.3	Ensure nosuid option set on /tmp partition	☐	☐
1.1.2.1.4	Ensure noexec option set on /tmp partition	☐	☐
1.1.2.2.1	Ensure /dev/shm is tmpfs or a separate partition	☐	☐
1.1.2.2.2	Ensure nodev option set on /dev/shm partition	☐	☐
1.1.2.2.3	Ensure nosuid option set on /dev/shm partition	☐	☐
1.1.2.2.4	Ensure noexec option set on /dev/shm partition	☐	☐
1.1.2.3.1	Ensure separate partition exists for /home	☐	☐
1.1.2.3.2	Ensure nodev option set on /home partition	☐	☐
1.1.2.3.3	Ensure nosuid option set on /home partition	☐	☐
1.1.2.4.1	Ensure separate partition exists for /var	☐	☐
1.1.2.4.2	Ensure nodev option set on /var partition	☐	☐
1.1.2.4.3	Ensure nosuid option set on /var partition	☐	☐
1.1.2.5.2	Ensure nodev option set on /var/tmp partition	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
1.1.2.5.3	Ensure nosuid option set on /var/tmp partition	☐	☐
1.1.2.5.4	Ensure noexec option set on /var/tmp partition	☐	☐
1.1.2.6.1	Ensure separate partition exists for /var/log	☐	☐
1.1.2.6.2	Ensure nodev option set on /var/log partition	☐	☐
1.1.2.6.3	Ensure nosuid option set on /var/log partition	☐	☐
1.1.2.6.4	Ensure noexec option set on /var/log partition	☐	☐
1.1.2.7.1	Ensure separate partition exists for /var/log/audit	☐	☐
1.1.2.7.2	Ensure nodev option set on /var/log/audit partition	☐	☐
1.1.2.7.3	Ensure nosuid option set on /var/log/audit partition	☐	☐
1.1.2.7.4	Ensure noexec option set on /var/log/audit partition	☐	☐
1.2.1.1	Ensure the source.list and .source files use the Signed-By option	☐	☐
1.2.1.2	Ensure weak dependencies are configured	☐	☐
1.2.1.3	Ensure access to gpg key files are configured	☐	☐
1.2.1.4	Ensure access to /etc/apt/trusted.gpg.d directory is configured	☐	☐
1.2.1.5	Ensure access to /etc/apt/auth.conf.d directory is configured	☐	☐
1.2.1.6	Ensure access to files in the /etc/apt/auth.conf.d/ directory is configured	☐	☐
1.2.1.7	Ensure access to /usr/share/keyrings directory is configured	☐	☐
1.2.1.8	Ensure access to /etc/apt/sources.list.d directory is configured	☐	☐
1.2.1.9	Ensure access to files in /etc/apt/sources.list.d are configured	☐	☐
1.2.1.10	Ensure repositories are set to https in /etc/apt/sources.list	☐	☐
1.2.1.11	Ensure repositories are set to https in /etc/apt/sources.list.d directory	☐	☐
1.2.1.12	Ensure AllowInsecureRepositories is configured	☐	☐
1.2.1.13	Ensure AllowWeakRepositories is configured	☐	☐
1.2.1.14	Ensure AllowDowngradeToInsecureRepositories is configured	☐	☐
1.2.1.15	Ensure Acquire::Check-Date is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
1.2.2.1	Ensure updates, patches, and additional security software are installed	☐	☐
1.3.1.1	Ensure apparmor packages are installed	☐	☐
1.3.1.2	Ensure AppArmor is enabled	☐	☐
1.3.1.3	Ensure all AppArmor Profiles are enforcing	☐	☐
1.3.1.4	Ensure kernel.apparmor_restrict_unprivileged_unconfined is enabled	☐	☐
1.4.1	Ensure bootloader password is set	☐	☐
1.4.2	Ensure access to bootloader config is configured	☐	☐
1.5.3	Ensure kernel.yama.ptrace_scope is configured	☐	☐
1.5.5	Ensure kernel.dmesg_restrict is configured	☐	☐
1.5.7	Ensure Automatic Error Reporting is configured	☐	☐
1.5.8	Ensure kernel.kptr_restrict is configured	☐	☐
1.5.9	Ensure kernel.randomize_va_space is configured	☐	☐
1.6.6	Ensure access to /etc/motd is configured	☐	☐
1.6.7	Ensure access to /etc/issue is configured	☐	☐
1.6.8	Ensure access to /etc/issue.net is configured	☐	☐
1.6.9	Ensure access to pam_motd file is configured	☐	☐
1.6.10	Ensure access to sshd warning banner is configured	☐	☐
1.6.11	Ensure update-notifier-motd.service is not in use	☐	☐
1.6.12	Ensure update-notifier-motd.timer is not in use	☐	☐
1.7.2	Ensure GDM disable-user-list is configured	☐	☐
1.7.3	Ensure GDM screen lock is configured	☐	☐
1.7.4	Ensure GDM automount is configured	☐	☐
1.7.5	Ensure GDM autorun-never is configured	☐	☐
1.7.6	Ensure XDMCP is not enabled	☐	☐
1.7.7	Ensure Xwayland is configured	☐	☐
2.1.1	Ensure autofs services are not in use	☐	☐
2.1.2	Ensure avahi daemon services are not in use	☐	☐
2.1.3	Ensure dhcp server services are not in use	☐	☐
2.1.4	Ensure dns server services are not in use	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
2.1.5	Ensure dnsmasq services are not in use	☐	☐
2.1.6	Ensure ftp server services are not in use	☐	☐
2.1.7	Ensure ldap server services are not in use	☐	☐
2.1.8	Ensure message access server services are not in use	☐	☐
2.1.9	Ensure network file system services are not in use	☐	☐
2.1.10	Ensure nis server services are not in use	☐	☐
2.1.11	Ensure print server services are not in use	☐	☐
2.1.12	Ensure rpcbind services are not in use	☐	☐
2.1.13	Ensure rsync services are not in use	☐	☐
2.1.14	Ensure samba file server services are not in use	☐	☐
2.1.15	Ensure snmp services are not in use	☐	☐
2.1.16	Ensure telnet server services are not in use	☐	☐
2.1.17	Ensure tftp server services are not in use	☐	☐
2.1.18	Ensure web proxy server services are not in use	☐	☐
2.1.19	Ensure web server services are not in use	☐	☐
2.1.20	Ensure xinetd services are not in use	☐	☐
2.1.21	Ensure X window server services are not in use	☐	☐
2.1.22	Ensure mail transfer agents are configured for local-only mode	☐	☐
2.1.23	Ensure only approved services are listening on a network interface	☐	☐
2.2.1	Ensure nis client is not installed	☐	☐
2.2.2	Ensure rsh client is not installed	☐	☐
2.2.3	Ensure talk client is not installed	☐	☐
2.2.4	Ensure telnet client is not installed	☐	☐
2.2.5	Ensure ldap client is not installed	☐	☐
2.2.6	Ensure ftp client is not installed	☐	☐
2.3.1.1	Ensure a single time synchronization daemon is in use	☐	☐
2.3.2.1	Ensure systemd-timesyncd configured with authorized timeserver	☐	☐
2.3.2.2	Ensure systemd-timesyncd is enabled and running	☐	☐
2.3.3.1	Ensure chrony is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
2.3.3.3	Ensure chrony is enabled and running	☐	☐
2.4.1.2	Ensure access to /etc/crontab is configured	☐	☐
2.4.1.3	Ensure access to /etc/cron.hourly is configured	☐	☐
2.4.1.4	Ensure access to /etc/cron.daily is configured	☐	☐
2.4.1.5	Ensure access to /etc/cron.weekly is configured	☐	☐
2.4.1.6	Ensure access to /etc/cron.monthly is configured	☐	☐
2.4.1.7	Ensure access to /etc/cron.yearly is configured	☐	☐
2.4.1.8	Ensure access to /etc/cron.d is configured	☐	☐
2.4.1.9	Ensure access to crontab is configured	☐	☐
2.4.2.1	Ensure access to at is configured	☐	☐
3.1.1	Ensure IPv6 status is identified	☐	☐
3.1.2	Ensure wireless interfaces are not available	☐	☐
3.1.3	Ensure bluetooth services are not in use	☐	☐
3.2.1	Ensure atm kernel module is not available	☐	☐
3.2.2	Ensure can kernel module is not available	☐	☐
3.2.3	Ensure dccp kernel module is not available	☐	☐
3.2.4	Ensure rds kernel module is not available	☐	☐
3.2.5	Ensure sctp kernel module is not available	☐	☐
3.2.6	Ensure tipc kernel module is not available	☐	☐
3.2.7	Ensure unneeded network protocol kernel modules are not available	☐	☐
3.3.1.1	Ensure net.ipv4.ip_forward is configured	☐	☐
3.3.1.2	Ensure net.ipv4.conf.all.forwarding is configured	☐	☐
3.3.1.3	Ensure net.ipv4.conf.default.forwarding is configured	☐	☐
3.3.1.4	Ensure net.ipv4.conf.all.send_redirects is configured	☐	☐
3.3.1.5	Ensure net.ipv4.conf.default.send_redirects is configured	☐	☐
3.3.1.6	Ensure net.ipv4.icmp_ignore_bogus_error_responses is configured	☐	☐
3.3.1.7	Ensure net.ipv4.icmp_echo_ignore_broadcasts is configured	☐	☐
3.3.1.8	Ensure net.ipv4.conf.all.accept_redirects is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
3.3.1.9	Ensure net.ipv4.conf.default.accept_redirects is configured	☐	☐
3.3.1.10	Ensure net.ipv4.conf.all.secure_redirects is configured	☐	☐
3.3.1.11	Ensure net.ipv4.conf.default.secure_redirects is configured	☐	☐
3.3.1.12	Ensure net.ipv4.conf.all.rp_filter is configured	☐	☐
3.3.1.13	Ensure net.ipv4.conf.default.rp_filter is configured	☐	☐
3.3.1.14	Ensure net.ipv4.conf.all.accept_source_route is configured	☐	☐
3.3.1.15	Ensure net.ipv4.conf.default.accept_source_route is configured	☐	☐
3.3.1.16	Ensure net.ipv4.conf.all.log_martians is configured	☐	☐
3.3.1.17	Ensure net.ipv4.conf.default.log_martians is configured	☐	☐
3.3.1.18	Ensure net.ipv4.tcp_syncookies is configured	☐	☐
3.3.1.19	Ensure net.ipv4.conf.all.route_localnet is configured	☐	☐
3.3.2.1	Ensure net.ipv6.conf.all.forwarding is configured	☐	☐
3.3.2.2	Ensure net.ipv6.conf.default.forwarding is configured	☐	☐
3.3.2.3	Ensure net.ipv6.conf.all.accept_redirects is configured	☐	☐
3.3.2.4	Ensure net.ipv6.conf.default.accept_redirects is configured	☐	☐
3.3.2.5	Ensure net.ipv6.conf.all.accept_source_route is configured	☐	☐
3.3.2.6	Ensure net.ipv6.conf.default.accept_source_route is configured	☐	☐
3.3.2.7	Ensure net.ipv6.conf.all.accept_ra is configured	☐	☐
3.3.2.8	Ensure net.ipv6.conf.default.accept_ra is configured	☐	☐
4.1.1	Ensure ufw is installed	☐	☐
4.1.2	Ensure ufw service is configured	☐	☐
4.1.3	Ensure ufw incoming default is configured	☐	☐
4.1.4	Ensure ufw outgoing default is configured	☐	☐
4.1.5	Ensure ufw routed default is configured	☐	☐
5.1.1	Ensure access to /etc/ssh/sshd_config is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
5.1.2	Ensure access to /etc/ssh/sshd_config.d/ directory is configured	☐	☐
5.1.3	Ensure access to SSH private host key files is configured	☐	☐
5.1.4	Ensure access to SSH public host key files is configured	☐	☐
5.1.5	Ensure sshd access is configured	☐	☐
5.1.7	Ensure sshd Ciphers are configured	☐	☐
5.1.9	Ensure sshd DisableForwarding is enabled	☐	☐
5.1.10	Ensure sshd GSSAPIAuthentication is disabled	☐	☐
5.1.11	Ensure sshd HostbasedAuthentication is disabled	☐	☐
5.1.12	Ensure sshd IgnoreRhosts is enabled	☐	☐
5.1.13	Ensure sshd KexAlgorithms is configured	☐	☐
5.1.14	Ensure sshd post-quantum cryptography key exchange algorithms are configured	☐	☐
5.1.16	Ensure sshd LogLevel is configured	☐	☐
5.1.17	Ensure sshd MACs are configured	☐	☐
5.1.18	Ensure sshd MaxAuthTries is configured	☐	☐
5.1.21	Ensure sshd PermitEmptyPasswords is disabled	☐	☐
5.1.22	Ensure sshd PermitRootLogin is disabled	☐	☐
5.1.24	Ensure sshd UsePAM is enabled	☐	☐
5.2.1	Ensure sudo is installed	☐	☐
5.2.2	Ensure sudo commands use pty	☐	☐
5.2.3	Ensure sudo log file exists	☐	☐
5.2.4	Ensure users must provide password for escalation	☐	☐
5.2.5	Ensure re-authentication for privilege escalation is not disabled globally	☐	☐
5.2.6	Ensure sudo timestamp_timeout is configured	☐	☐
5.2.7	Ensure access to the su command is restricted	☐	☐
5.3.1.3	Ensure latest version of libpam-pwquality is installed	☐	☐
5.3.1.4	Ensure latest version of cracklib-runtime is installed	☐	☐
5.3.2.1	Ensure pam_unix module is enabled	☐	☐
5.3.2.2	Ensure pam_faillock module is enabled	☐	☐
5.3.2.3	Ensure pam_pwquality module is enabled	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
5.3.2.4	Ensure pam_pwhistory module is enabled	☐	☐
5.3.3.1.1	Ensure password failed attempts lockout is configured	☐	☐
5.3.3.1.2	Ensure password unlock time is configured	☐	☐
5.3.3.1.3	Ensure password failed attempts lockout includes root account	☐	☐
5.3.3.2.1	Ensure password number of changed characters is configured	☐	☐
5.3.3.2.2	Ensure password length is configured	☐	☐
5.3.3.2.3	Ensure password complexity is configured	☐	☐
5.3.3.2.4	Ensure password same consecutive characters is configured	☐	☐
5.3.3.2.5	Ensure password maximum sequential characters is configured	☐	☐
5.3.3.2.6	Ensure password dictionary check is enabled	☐	☐
5.3.3.2.7	Ensure password quality checking is enforced	☐	☐
5.3.3.2.8	Ensure password quality is enforced for the root user	☐	☐
5.3.3.3.1	Ensure password history remember is configured	☐	☐
5.3.3.3.2	Ensure password history is enforced for the root user	☐	☐
5.3.3.3.3	Ensure pam_pwhistory includes use_authtok	☐	☐
5.3.3.4.1	Ensure pam_unix does not include nullok	☐	☐
5.3.3.4.2	Ensure pam_unix does not include remember	☐	☐
5.3.3.4.3	Ensure pam_unix includes a strong password hashing algorithm	☐	☐
5.3.3.4.4	Ensure pam_unix includes use_authtok	☐	☐
5.4.1.1	Ensure password expiration is configured	☐	☐
5.4.1.2	Ensure minimum password days is configured	☐	☐
5.4.1.3	Ensure password expiration warning days is configured	☐	☐
5.4.1.4	Ensure strong password hashing algorithm is configured	☐	☐
5.4.1.5	Ensure inactive password lock is configured	☐	☐
5.4.1.6	Ensure all users last password change date is in the past	☐	☐
5.4.2.1	Ensure root is the only UID 0 account	☐	☐
5.4.2.2	Ensure root is the only GID 0 account	☐	☐
5.4.2.3	Ensure group root is the only GID 0 group	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
5.4.2.4	Ensure root account access is controlled	☐	☐
5.4.2.6	Ensure root user umask is configured	☐	☐
5.4.2.7	Ensure system accounts do not have a valid login shell	☐	☐
5.4.2.8	Ensure accounts without a valid login shell are locked	☐	☐
5.4.3.2	Ensure default user shell timeout is configured	☐	☐
5.4.3.3	Ensure default user umask is configured	☐	☐
6.1.1.1.1	Ensure journald service is active	☐	☐
6.1.1.1.2	Ensure systemd-journal-remote service is not in use	☐	☐
6.1.1.1.3	Ensure journald is configured to send logs to rsyslog	☐	☐
6.1.1.1.4	Ensure journald log file access is configured	☐	☐
6.1.1.1.5	Ensure journald log file rotation is configured	☐	☐
6.1.1.1.6	Ensure journald Storage is configured	☐	☐
6.1.1.1.7	Ensure journald Compress is configured	☐	☐
6.1.2.1	Ensure rsyslog is installed	☐	☐
6.1.2.2	Ensure rsyslog service is enabled and active	☐	☐
6.1.2.3	Ensure rsyslog log file creation mode is configured	☐	☐
6.1.2.4	Ensure rsyslog logging is configured	☐	☐
6.1.2.5	Ensure rsyslog is configured to send logs to a remote log host	☐	☐
6.1.2.6	Ensure rsyslog is not configured to receive logs from a remote client	☐	☐
6.1.2.7	Ensure logrotate is configured	☐	☐
6.1.2.8	Ensure rsyslog-gnutls is installed	☐	☐
6.1.2.9	Ensure rsyslog forwarding uses gtls	☐	☐
6.1.2.10	Ensure rsyslog CA certificates are configured	☐	☐
6.1.3.1	Ensure access to all logfiles has been configured	☐	☐
6.2.1.1	Ensure auditd packages are installed	☐	☐
6.2.1.2	Ensure auditd service is enabled and active	☐	☐
6.2.1.3	Ensure auditing for processes that start prior to auditd is enabled	☐	☐
6.2.1.4	Ensure audit_backlog_limit is configured	☐	☐
6.2.2.1	Ensure audit log storage size is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
6.2.2.2	Ensure audit logs are not automatically deleted	☐	☐
6.2.2.3	Ensure system is disabled when audit logs are full	☐	☐
6.2.2.4	Ensure system warns when audit logs are low on space	☐	☐
6.2.3.1	Ensure modification of the /etc/sudoers file is collected	☐	☐
6.2.3.2	Ensure actions as another user are always logged	☐	☐
6.2.3.3	Ensure events that modify the sudo log file are collected	☐	☐
6.2.3.4	Ensure events that modify date and time information are collected	☐	☐
6.2.3.5	Ensure events that modify sethostname and setdomainname are collected	☐	☐
6.2.3.6	Ensure events that modify /etc/issue and /etc/issue.net are collected	☐	☐
6.2.3.7	Ensure events that modify /etc/hosts and /etc/hostname are collected	☐	☐
6.2.3.8	Ensure events that modify the system's network environment are collected	☐	☐
6.2.3.9	Ensure events that modify /etc/NetworkManager directory are collected	☐	☐
6.2.3.10	Ensure use of privileged commands are collected	☐	☐
6.2.3.11	Ensure unsuccessful file access attempts are collected	☐	☐
6.2.3.12	Ensure events that modify /etc/group information are collected	☐	☐
6.2.3.13	Ensure events that modify /etc/passwd information are collected	☐	☐
6.2.3.14	Ensure events that modify /etc/shadow and /etc/gshadow are collected	☐	☐
6.2.3.15	Ensure events that modify /etc/security/opasswd are collected	☐	☐
6.2.3.16	Ensure events that modify /etc/nsswitch.conf file are collected	☐	☐
6.2.3.17	Ensure events that modify /etc/pam.conf and /etc/pam.d/ information are collected	☐	☐
6.2.3.18	Ensure discretionary access control permission modification events chmod,fchmod,fchmodat,fchmodat2 are collected	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
6.2.3.19	Ensure discretionary access control permission modification events chown,fchown,lchown,fchownat are collected	☐	☐
6.2.3.20	Ensure discretionary access control permission modification events setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr collected	☐	☐
6.2.3.21	Ensure successful file system mounts are collected	☐	☐
6.2.3.22	Ensure session initiation information is collected	☐	☐
6.2.3.23	Ensure login and logout events are collected	☐	☐
6.2.3.24	Ensure unlink file deletion events by users are collected	☐	☐
6.2.3.25	Ensure rename file deletion events by users are collected	☐	☐
6.2.3.26	Ensure events that modify the system's Mandatory Access Controls are collected	☐	☐
6.2.3.27	Ensure successful and unsuccessful attempts to use the chcon command are collected	☐	☐
6.2.3.28	Ensure successful and unsuccessful attempts to use the setfacl command are collected	☐	☐
6.2.3.29	Ensure successful and unsuccessful attempts to use the chacl command are collected	☐	☐
6.2.3.30	Ensure successful and unsuccessful attempts to use the usermod command are collected	☐	☐
6.2.3.31	Ensure kernel module loading unloading and modification is collected	☐	☐
6.2.3.32	Ensure kernel "init_module" and "finit_module" loading unloading and modification is collected	☐	☐
6.2.3.33	Ensure kernel "delete_module" loading unloading and modification is collected	☐	☐
6.2.3.34	Ensure the audit configuration is loaded regardless of errors	☐	☐
6.2.3.35	Ensure the audit configuration is immutable	☐	☐
6.2.3.36	Ensure the running and on disk configuration is the same	☐	☐
6.2.4.1	Ensure audit log files mode is configured	☐	☐
6.2.4.2	Ensure audit log files owner is configured	☐	☐
6.2.4.3	Ensure audit log files group owner is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
6.2.4.4	Ensure the audit log file directory mode is configured	☐	☐
6.2.4.5	Ensure audit configuration files mode is configured	☐	☐
6.2.4.6	Ensure audit configuration files owner is configured	☐	☐
6.2.4.7	Ensure audit configuration files group owner is configured	☐	☐
6.2.4.8	Ensure audit tools mode is configured	☐	☐
6.2.4.9	Ensure audit tools owner is configured	☐	☐
6.2.4.10	Ensure audit tools group owner is configured	☐	☐
7.1.1	Ensure access to /etc/passwd is configured	☐	☐
7.1.2	Ensure access to /etc/passwd- is configured	☐	☐
7.1.3	Ensure access to /etc/group is configured	☐	☐
7.1.4	Ensure access to /etc/group- is configured	☐	☐
7.1.5	Ensure access to /etc/shadow is configured	☐	☐
7.1.6	Ensure access to /etc/shadow- is configured	☐	☐
7.1.7	Ensure access to /etc/gshadow is configured	☐	☐
7.1.8	Ensure access to /etc/gshadow- is configured	☐	☐
7.1.9	Ensure access to /etc/shells is configured	☐	☐
7.1.10	Ensure access to /etc/security/opasswd is configured	☐	☐
7.1.11	Ensure world writable files and directories are secured	☐	☐
7.1.12	Ensure no files or directories without an owner and a group exist	☐	☐
7.1.13	Ensure SUID and SGID files are reviewed	☐	☐
7.2.1	Ensure accounts in /etc/passwd use shadowed passwords	☐	☐
7.2.2	Ensure /etc/shadow password fields are not empty	☐	☐
7.2.3	Ensure all groups in /etc/passwd exist in /etc/group	☐	☐
7.2.4	Ensure shadow group is empty	☐	☐
7.2.9	Ensure local interactive user home directories are configured	☐	☐
7.2.10	Ensure local interactive user dot files access is configured	☐	☐
7.2.11	Ensure local interactive user dot directories access is configured	☐	☐

Internal Only - General

Internal Only - General

Appendix: CIS Controls v8 IG 3 Mapped Recommendations

Recommendation		Set Correctly	
		Yes	No
1.1.1.1	Ensure cramfs kernel module is not available	☐	☐
1.1.1.2	Ensure freevxfs kernel module is not available	☐	☐
1.1.1.3	Ensure hfs kernel module is not available	☐	☐
1.1.1.4	Ensure hfsplus kernel module is not available	☐	☐
1.1.1.5	Ensure jffs2 kernel module is not available	☐	☐
1.1.1.6	Ensure overlay kernel module is not available	☐	☐
1.1.1.7	Ensure squashfs kernel module is not available	☐	☐
1.1.1.8	Ensure udf kernel module is not available	☐	☐
1.1.1.9	Ensure firewire-core kernel module is not available	☐	☐
1.1.1.10	Ensure usb-storage kernel module is not available	☐	☐
1.1.1.11	Ensure unused filesystems kernel modules are not available	☐	☐
1.1.2.1.1	Ensure /tmp is tmpfs or a separate partition	☐	☐
1.1.2.1.2	Ensure nodev option set on /tmp partition	☐	☐
1.1.2.1.3	Ensure nosuid option set on /tmp partition	☐	☐
1.1.2.1.4	Ensure noexec option set on /tmp partition	☐	☐
1.1.2.2.1	Ensure /dev/shm is tmpfs or a separate partition	☐	☐
1.1.2.2.2	Ensure nodev option set on /dev/shm partition	☐	☐
1.1.2.2.3	Ensure nosuid option set on /dev/shm partition	☐	☐
1.1.2.2.4	Ensure noexec option set on /dev/shm partition	☐	☐
1.1.2.3.1	Ensure separate partition exists for /home	☐	☐
1.1.2.3.2	Ensure nodev option set on /home partition	☐	☐
1.1.2.3.3	Ensure nosuid option set on /home partition	☐	☐
1.1.2.4.1	Ensure separate partition exists for /var	☐	☐
1.1.2.4.2	Ensure nodev option set on /var partition	☐	☐
1.1.2.4.3	Ensure nosuid option set on /var partition	☐	☐
1.1.2.5.2	Ensure nodev option set on /var/tmp partition	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
1.1.2.5.3	Ensure nosuid option set on /var/tmp partition	☐	☐
1.1.2.5.4	Ensure noexec option set on /var/tmp partition	☐	☐
1.1.2.6.1	Ensure separate partition exists for /var/log	☐	☐
1.1.2.6.2	Ensure nodev option set on /var/log partition	☐	☐
1.1.2.6.3	Ensure nosuid option set on /var/log partition	☐	☐
1.1.2.6.4	Ensure noexec option set on /var/log partition	☐	☐
1.1.2.7.1	Ensure separate partition exists for /var/log/audit	☐	☐
1.1.2.7.2	Ensure nodev option set on /var/log/audit partition	☐	☐
1.1.2.7.3	Ensure nosuid option set on /var/log/audit partition	☐	☐
1.1.2.7.4	Ensure noexec option set on /var/log/audit partition	☐	☐
1.2.1.1	Ensure the source.list and .source files use the Signed-By option	☐	☐
1.2.1.2	Ensure weak dependencies are configured	☐	☐
1.2.1.3	Ensure access to gpg key files are configured	☐	☐
1.2.1.4	Ensure access to /etc/apt/trusted.gpg.d directory is configured	☐	☐
1.2.1.5	Ensure access to /etc/apt/auth.conf.d directory is configured	☐	☐
1.2.1.6	Ensure access to files in the /etc/apt/auth.conf.d/ directory is configured	☐	☐
1.2.1.7	Ensure access to /usr/share/keyrings directory is configured	☐	☐
1.2.1.8	Ensure access to /etc/apt/sources.list.d directory is configured	☐	☐
1.2.1.9	Ensure access to files in /etc/apt/sources.list.d are configured	☐	☐
1.2.1.10	Ensure repositories are set to https in /etc/apt/sources.list	☐	☐
1.2.1.11	Ensure repositories are set to https in /etc/apt/sources.list.d directory	☐	☐
1.2.1.12	Ensure AllowInsecureRepositories is configured	☐	☐
1.2.1.13	Ensure AllowWeakRepositories is configured	☐	☐
1.2.1.14	Ensure AllowDowngradeToInsecureRepositories is configured	☐	☐
1.2.1.15	Ensure Acquire::Check-Date is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
1.2.2.1	Ensure updates, patches, and additional security software are installed	☐	☐
1.3.1.1	Ensure apparmor packages are installed	☐	☐
1.3.1.2	Ensure AppArmor is enabled	☐	☐
1.3.1.3	Ensure all AppArmor Profiles are enforcing	☐	☐
1.3.1.4	Ensure kernel.apparmor_restrict_unprivileged_unconfined is enabled	☐	☐
1.4.1	Ensure bootloader password is set	☐	☐
1.4.2	Ensure access to bootloader config is configured	☐	☐
1.5.3	Ensure kernel.yama.ptrace_scope is configured	☐	☐
1.5.5	Ensure kernel.dmesg_restrict is configured	☐	☐
1.5.6	Ensure prelink is not installed	☐	☐
1.5.7	Ensure Automatic Error Reporting is configured	☐	☐
1.5.8	Ensure kernel.kptr_restrict is configured	☐	☐
1.5.9	Ensure kernel.randomize_va_space is configured	☐	☐
1.6.6	Ensure access to /etc/motd is configured	☐	☐
1.6.7	Ensure access to /etc/issue is configured	☐	☐
1.6.8	Ensure access to /etc/issue.net is configured	☐	☐
1.6.9	Ensure access to pam_motd file is configured	☐	☐
1.6.10	Ensure access to sshd warning banner is configured	☐	☐
1.6.11	Ensure update-notifier-motd.service is not in use	☐	☐
1.6.12	Ensure update-notifier-motd.timer is not in use	☐	☐
1.7.2	Ensure GDM disable-user-list is configured	☐	☐
1.7.3	Ensure GDM screen lock is configured	☐	☐
1.7.4	Ensure GDM automount is configured	☐	☐
1.7.5	Ensure GDM autorun-never is configured	☐	☐
1.7.6	Ensure XDMCP is not enabled	☐	☐
1.7.7	Ensure Xwayland is configured	☐	☐
2.1.1	Ensure autofs services are not in use	☐	☐
2.1.2	Ensure avahi daemon services are not in use	☐	☐
2.1.3	Ensure dhcp server services are not in use	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
2.1.4	Ensure dns server services are not in use	☐	☐
2.1.5	Ensure dnsmasq services are not in use	☐	☐
2.1.6	Ensure ftp server services are not in use	☐	☐
2.1.7	Ensure ldap server services are not in use	☐	☐
2.1.8	Ensure message access server services are not in use	☐	☐
2.1.9	Ensure network file system services are not in use	☐	☐
2.1.10	Ensure nis server services are not in use	☐	☐
2.1.11	Ensure print server services are not in use	☐	☐
2.1.12	Ensure rpcbind services are not in use	☐	☐
2.1.13	Ensure rsync services are not in use	☐	☐
2.1.14	Ensure samba file server services are not in use	☐	☐
2.1.15	Ensure snmp services are not in use	☐	☐
2.1.16	Ensure telnet server services are not in use	☐	☐
2.1.17	Ensure tftp server services are not in use	☐	☐
2.1.18	Ensure web proxy server services are not in use	☐	☐
2.1.19	Ensure web server services are not in use	☐	☐
2.1.20	Ensure xinetd services are not in use	☐	☐
2.1.21	Ensure X window server services are not in use	☐	☐
2.1.22	Ensure mail transfer agents are configured for local-only mode	☐	☐
2.1.23	Ensure only approved services are listening on a network interface	☐	☐
2.2.1	Ensure nis client is not installed	☐	☐
2.2.2	Ensure rsh client is not installed	☐	☐
2.2.3	Ensure talk client is not installed	☐	☐
2.2.4	Ensure telnet client is not installed	☐	☐
2.2.5	Ensure ldap client is not installed	☐	☐
2.2.6	Ensure ftp client is not installed	☐	☐
2.3.1.1	Ensure a single time synchronization daemon is in use	☐	☐
2.3.2.1	Ensure systemd-timesyncd configured with authorized timeserver	☐	☐
2.3.2.2	Ensure systemd-timesyncd is enabled and running	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
2.3.3.1	Ensure chrony is configured	☐	☐
2.3.3.3	Ensure chrony is enabled and running	☐	☐
2.4.1.2	Ensure access to /etc/crontab is configured	☐	☐
2.4.1.3	Ensure access to /etc/cron.hourly is configured	☐	☐
2.4.1.4	Ensure access to /etc/cron.daily is configured	☐	☐
2.4.1.5	Ensure access to /etc/cron.weekly is configured	☐	☐
2.4.1.6	Ensure access to /etc/cron.monthly is configured	☐	☐
2.4.1.7	Ensure access to /etc/cron.yearly is configured	☐	☐
2.4.1.8	Ensure access to /etc/cron.d is configured	☐	☐
2.4.1.9	Ensure access to crontab is configured	☐	☐
2.4.2.1	Ensure access to at is configured	☐	☐
3.1.1	Ensure IPv6 status is identified	☐	☐
3.1.2	Ensure wireless interfaces are not available	☐	☐
3.1.3	Ensure bluetooth services are not in use	☐	☐
3.2.1	Ensure atm kernel module is not available	☐	☐
3.2.2	Ensure can kernel module is not available	☐	☐
3.2.3	Ensure dccp kernel module is not available	☐	☐
3.2.4	Ensure rds kernel module is not available	☐	☐
3.2.5	Ensure sctp kernel module is not available	☐	☐
3.2.6	Ensure tipc kernel module is not available	☐	☐
3.2.7	Ensure unneeded network protocol kernel modules are not available	☐	☐
3.3.1.1	Ensure net.ipv4.ip_forward is configured	☐	☐
3.3.1.2	Ensure net.ipv4.conf.all.forwarding is configured	☐	☐
3.3.1.3	Ensure net.ipv4.conf.default.forwarding is configured	☐	☐
3.3.1.4	Ensure net.ipv4.conf.all.send_redirects is configured	☐	☐
3.3.1.5	Ensure net.ipv4.conf.default.send_redirects is configured	☐	☐
3.3.1.6	Ensure net.ipv4.icmp_ignore_bogus_error_responses is configured	☐	☐
3.3.1.7	Ensure net.ipv4.icmp_echo_ignore_broadcasts is configured	☐	☐
3.3.1.8	Ensure net.ipv4.conf.all.accept_redirects is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
3.3.1.9	Ensure net.ipv4.conf.default.accept_redirects is configured	☐	☐
3.3.1.10	Ensure net.ipv4.conf.all.secure_redirects is configured	☐	☐
3.3.1.11	Ensure net.ipv4.conf.default.secure_redirects is configured	☐	☐
3.3.1.12	Ensure net.ipv4.conf.all.rp_filter is configured	☐	☐
3.3.1.13	Ensure net.ipv4.conf.default.rp_filter is configured	☐	☐
3.3.1.14	Ensure net.ipv4.conf.all.accept_source_route is configured	☐	☐
3.3.1.15	Ensure net.ipv4.conf.default.accept_source_route is configured	☐	☐
3.3.1.16	Ensure net.ipv4.conf.all.log_martians is configured	☐	☐
3.3.1.17	Ensure net.ipv4.conf.default.log_martians is configured	☐	☐
3.3.1.18	Ensure net.ipv4.tcp_syncookies is configured	☐	☐
3.3.1.19	Ensure net.ipv4.conf.all.route_localnet is configured	☐	☐
3.3.2.1	Ensure net.ipv6.conf.all.forwarding is configured	☐	☐
3.3.2.2	Ensure net.ipv6.conf.default.forwarding is configured	☐	☐
3.3.2.3	Ensure net.ipv6.conf.all.accept_redirects is configured	☐	☐
3.3.2.4	Ensure net.ipv6.conf.default.accept_redirects is configured	☐	☐
3.3.2.5	Ensure net.ipv6.conf.all.accept_source_route is configured	☐	☐
3.3.2.6	Ensure net.ipv6.conf.default.accept_source_route is configured	☐	☐
3.3.2.7	Ensure net.ipv6.conf.all.accept_ra is configured	☐	☐
3.3.2.8	Ensure net.ipv6.conf.default.accept_ra is configured	☐	☐
4.1.1	Ensure ufw is installed	☐	☐
4.1.2	Ensure ufw service is configured	☐	☐
4.1.3	Ensure ufw incoming default is configured	☐	☐
4.1.4	Ensure ufw outgoing default is configured	☐	☐
4.1.5	Ensure ufw routed default is configured	☐	☐
5.1.1	Ensure access to /etc/ssh/sshd_config is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
5.1.2	Ensure access to /etc/ssh/sshd_config.d/ directory is configured	☐	☐
5.1.3	Ensure access to SSH private host key files is configured	☐	☐
5.1.4	Ensure access to SSH public host key files is configured	☐	☐
5.1.5	Ensure sshd access is configured	☐	☐
5.1.7	Ensure sshd Ciphers are configured	☐	☐
5.1.9	Ensure sshd DisableForwarding is enabled	☐	☐
5.1.10	Ensure sshd GSSAPIAuthentication is disabled	☐	☐
5.1.11	Ensure sshd HostbasedAuthentication is disabled	☐	☐
5.1.12	Ensure sshd IgnoreRhosts is enabled	☐	☐
5.1.13	Ensure sshd KexAlgorithms is configured	☐	☐
5.1.14	Ensure sshd post-quantum cryptography key exchange algorithms are configured	☐	☐
5.1.16	Ensure sshd LogLevel is configured	☐	☐
5.1.17	Ensure sshd MACs are configured	☐	☐
5.1.18	Ensure sshd MaxAuthTries is configured	☐	☐
5.1.21	Ensure sshd PermitEmptyPasswords is disabled	☐	☐
5.1.22	Ensure sshd PermitRootLogin is disabled	☐	☐
5.1.24	Ensure sshd UsePAM is enabled	☐	☐
5.2.1	Ensure sudo is installed	☐	☐
5.2.2	Ensure sudo commands use pty	☐	☐
5.2.3	Ensure sudo log file exists	☐	☐
5.2.4	Ensure users must provide password for escalation	☐	☐
5.2.5	Ensure re-authentication for privilege escalation is not disabled globally	☐	☐
5.2.6	Ensure sudo timestamp_timeout is configured	☐	☐
5.2.7	Ensure access to the su command is restricted	☐	☐
5.3.1.3	Ensure latest version of libpam-pwquality is installed	☐	☐
5.3.1.4	Ensure latest version of cracklib-runtime is installed	☐	☐
5.3.2.1	Ensure pam_unix module is enabled	☐	☐
5.3.2.2	Ensure pam_faillock module is enabled	☐	☐
5.3.2.3	Ensure pam_pwquality module is enabled	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
5.3.2.4	Ensure pam_pwhistory module is enabled	☐	☐
5.3.3.1.1	Ensure password failed attempts lockout is configured	☐	☐
5.3.3.1.2	Ensure password unlock time is configured	☐	☐
5.3.3.1.3	Ensure password failed attempts lockout includes root account	☐	☐
5.3.3.2.1	Ensure password number of changed characters is configured	☐	☐
5.3.3.2.2	Ensure password length is configured	☐	☐
5.3.3.2.3	Ensure password complexity is configured	☐	☐
5.3.3.2.4	Ensure password same consecutive characters is configured	☐	☐
5.3.3.2.5	Ensure password maximum sequential characters is configured	☐	☐
5.3.3.2.6	Ensure password dictionary check is enabled	☐	☐
5.3.3.2.7	Ensure password quality checking is enforced	☐	☐
5.3.3.2.8	Ensure password quality is enforced for the root user	☐	☐
5.3.3.3.1	Ensure password history remember is configured	☐	☐
5.3.3.3.2	Ensure password history is enforced for the root user	☐	☐
5.3.3.3.3	Ensure pam_pwhistory includes use_authtok	☐	☐
5.3.3.4.1	Ensure pam_unix does not include nullok	☐	☐
5.3.3.4.2	Ensure pam_unix does not include remember	☐	☐
5.3.3.4.3	Ensure pam_unix includes a strong password hashing algorithm	☐	☐
5.3.3.4.4	Ensure pam_unix includes use_authtok	☐	☐
5.4.1.1	Ensure password expiration is configured	☐	☐
5.4.1.2	Ensure minimum password days is configured	☐	☐
5.4.1.3	Ensure password expiration warning days is configured	☐	☐
5.4.1.4	Ensure strong password hashing algorithm is configured	☐	☐
5.4.1.5	Ensure inactive password lock is configured	☐	☐
5.4.1.6	Ensure all users last password change date is in the past	☐	☐
5.4.2.1	Ensure root is the only UID 0 account	☐	☐
5.4.2.2	Ensure root is the only GID 0 account	☐	☐
5.4.2.3	Ensure group root is the only GID 0 group	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
5.4.2.4	Ensure root account access is controlled	☐	☐
5.4.2.6	Ensure root user umask is configured	☐	☐
5.4.2.7	Ensure system accounts do not have a valid login shell	☐	☐
5.4.2.8	Ensure accounts without a valid login shell are locked	☐	☐
5.4.3.2	Ensure default user shell timeout is configured	☐	☐
5.4.3.3	Ensure default user umask is configured	☐	☐
6.1.1.1.1	Ensure journald service is active	☐	☐
6.1.1.1.2	Ensure systemd-journal-remote service is not in use	☐	☐
6.1.1.1.3	Ensure journald is configured to send logs to rsyslog	☐	☐
6.1.1.1.4	Ensure journald log file access is configured	☐	☐
6.1.1.1.5	Ensure journald log file rotation is configured	☐	☐
6.1.1.1.6	Ensure journald Storage is configured	☐	☐
6.1.1.1.7	Ensure journald Compress is configured	☐	☐
6.1.2.1	Ensure rsyslog is installed	☐	☐
6.1.2.2	Ensure rsyslog service is enabled and active	☐	☐
6.1.2.3	Ensure rsyslog log file creation mode is configured	☐	☐
6.1.2.4	Ensure rsyslog logging is configured	☐	☐
6.1.2.5	Ensure rsyslog is configured to send logs to a remote log host	☐	☐
6.1.2.6	Ensure rsyslog is not configured to receive logs from a remote client	☐	☐
6.1.2.7	Ensure logrotate is configured	☐	☐
6.1.2.8	Ensure rsyslog-gnutls is installed	☐	☐
6.1.2.9	Ensure rsyslog forwarding uses gtls	☐	☐
6.1.2.10	Ensure rsyslog CA certificates are configured	☐	☐
6.1.3.1	Ensure access to all logfiles has been configured	☐	☐
6.2.1.1	Ensure auditd packages are installed	☐	☐
6.2.1.2	Ensure auditd service is enabled and active	☐	☐
6.2.1.3	Ensure auditing for processes that start prior to auditd is enabled	☐	☐
6.2.1.4	Ensure audit_backlog_limit is configured	☐	☐
6.2.2.1	Ensure audit log storage size is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
6.2.2.2	Ensure audit logs are not automatically deleted	☐	☐
6.2.2.3	Ensure system is disabled when audit logs are full	☐	☐
6.2.2.4	Ensure system warns when audit logs are low on space	☐	☐
6.2.3.1	Ensure modification of the /etc/sudoers file is collected	☐	☐
6.2.3.2	Ensure actions as another user are always logged	☐	☐
6.2.3.3	Ensure events that modify the sudo log file are collected	☐	☐
6.2.3.4	Ensure events that modify date and time information are collected	☐	☐
6.2.3.5	Ensure events that modify sethostname and setdomainname are collected	☐	☐
6.2.3.6	Ensure events that modify /etc/issue and /etc/issue.net are collected	☐	☐
6.2.3.7	Ensure events that modify /etc/hosts and /etc/hostname are collected	☐	☐
6.2.3.8	Ensure events that modify the system's network environment are collected	☐	☐
6.2.3.9	Ensure events that modify /etc/NetworkManager directory are collected	☐	☐
6.2.3.10	Ensure use of privileged commands are collected	☐	☐
6.2.3.11	Ensure unsuccessful file access attempts are collected	☐	☐
6.2.3.12	Ensure events that modify /etc/group information are collected	☐	☐
6.2.3.13	Ensure events that modify /etc/passwd information are collected	☐	☐
6.2.3.14	Ensure events that modify /etc/shadow and /etc/gshadow are collected	☐	☐
6.2.3.15	Ensure events that modify /etc/security/opasswd are collected	☐	☐
6.2.3.16	Ensure events that modify /etc/nsswitch.conf file are collected	☐	☐
6.2.3.17	Ensure events that modify /etc/pam.conf and /etc/pam.d/ information are collected	☐	☐
6.2.3.18	Ensure discretionary access control permission modification events chmod,fchmod,fchmodat,fchmodat2 are collected	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
6.2.3.19	Ensure discretionary access control permission modification events chown,fchown,lchown,fchownat are collected	☐	☐
6.2.3.20	Ensure discretionary access control permission modification events setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr collected	☐	☐
6.2.3.21	Ensure successful file system mounts are collected	☐	☐
6.2.3.22	Ensure session initiation information is collected	☐	☐
6.2.3.23	Ensure login and logout events are collected	☐	☐
6.2.3.24	Ensure unlink file deletion events by users are collected	☐	☐
6.2.3.25	Ensure rename file deletion events by users are collected	☐	☐
6.2.3.26	Ensure events that modify the system's Mandatory Access Controls are collected	☐	☐
6.2.3.27	Ensure successful and unsuccessful attempts to use the chcon command are collected	☐	☐
6.2.3.28	Ensure successful and unsuccessful attempts to use the setfacl command are collected	☐	☐
6.2.3.29	Ensure successful and unsuccessful attempts to use the chacl command are collected	☐	☐
6.2.3.30	Ensure successful and unsuccessful attempts to use the usermod command are collected	☐	☐
6.2.3.31	Ensure kernel module loading unloading and modification is collected	☐	☐
6.2.3.32	Ensure kernel "init_module" and "finit_module" loading unloading and modification is collected	☐	☐
6.2.3.33	Ensure kernel "delete_module" loading unloading and modification is collected	☐	☐
6.2.3.34	Ensure the audit configuration is loaded regardless of errors	☐	☐
6.2.3.35	Ensure the audit configuration is immutable	☐	☐
6.2.3.36	Ensure the running and on disk configuration is the same	☐	☐
6.2.4.1	Ensure audit log files mode is configured	☐	☐
6.2.4.2	Ensure audit log files owner is configured	☐	☐
6.2.4.3	Ensure audit log files group owner is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
6.2.4.4	Ensure the audit log file directory mode is configured	☐	☐
6.2.4.5	Ensure audit configuration files mode is configured	☐	☐
6.2.4.6	Ensure audit configuration files owner is configured	☐	☐
6.2.4.7	Ensure audit configuration files group owner is configured	☐	☐
6.2.4.8	Ensure audit tools mode is configured	☐	☐
6.2.4.9	Ensure audit tools owner is configured	☐	☐
6.2.4.10	Ensure audit tools group owner is configured	☐	☐
6.3.1	Ensure AIDE is installed	☐	☐
6.3.2	Ensure filesystem integrity is regularly checked	☐	☐
7.1.1	Ensure access to /etc/passwd is configured	☐	☐
7.1.2	Ensure access to /etc/passwd- is configured	☐	☐
7.1.3	Ensure access to /etc/group is configured	☐	☐
7.1.4	Ensure access to /etc/group- is configured	☐	☐
7.1.5	Ensure access to /etc/shadow is configured	☐	☐
7.1.6	Ensure access to /etc/shadow- is configured	☐	☐
7.1.7	Ensure access to /etc/gshadow is configured	☐	☐
7.1.8	Ensure access to /etc/gshadow- is configured	☐	☐
7.1.9	Ensure access to /etc/shells is configured	☐	☐
7.1.10	Ensure access to /etc/security/opasswd is configured	☐	☐
7.1.11	Ensure world writable files and directories are secured	☐	☐
7.1.12	Ensure no files or directories without an owner and a group exist	☐	☐
7.1.13	Ensure SUID and SGID files are reviewed	☐	☐
7.2.1	Ensure accounts in /etc/passwd use shadowed passwords	☐	☐
7.2.2	Ensure /etc/shadow password fields are not empty	☐	☐
7.2.3	Ensure all groups in /etc/passwd exist in /etc/group	☐	☐
7.2.4	Ensure shadow group is empty	☐	☐
7.2.9	Ensure local interactive user home directories are configured	☐	☐
7.2.10	Ensure local interactive user dot files access is configured	☐	☐

Internal Only - General

Recommendation		Set Correctly	
		Yes	No
7.2.11	Ensure local interactive user dot directories access is configured	☐	☐

Internal Only - General

Appendix: CIS Controls v8 Unmapped Recommendations

Recommendation		Set Correctly	
		Yes	No
	No unmapped recommendations to CIS Controls v8	☐	☐

Internal Only - General

Appendix: Change History

Date	Version	Changes for this version

Internal Only - General