CIS AlmaLinux OS 9 Benchmark

v3.0.0 - 08-28-2026

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.

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) 28

1.1.1.3 Ensure hfs kernel module is not available (Automated) 32
1.1.1.4 Ensure hfsplus kernel module is not available (Automated) 35
1.1.1.5 Ensure jffs2 kernel module is not available (Automated) 39
1.1.1.6 Ensure overlay kernel module is not available (Automated) 43
1.1.1.7 Ensure squashfs kernel module is not available (Automated) 47
1.1.1.8 Ensure udf kernel module is not available (Automated) 51
1.1.1.9 Ensure firewire-core kernel module is not available (Automated) 54
1.1.1.10 Ensure usb-storage kernel module is not available (Automated) 58
1.1.1.11 Ensure unused filesystems kernel modules are not available (Manual) 62
1.1.2 Configure Filesystem Partitions 69
1.1.2.1 Configure /tmp 71
1.1.2.1.1 Ensure /tmp is tmpfs or a separate partition (Automated) 72
1.1.2.1.2 Ensure nodev option set on /tmp partition (Automated) 76
1.1.2.1.3 Ensure nosuid option set on /tmp partition (Automated) 78
1.1.2.1.4 Ensure noexec option set on /tmp partition (Automated) 80
1.1.2.2 Configure /dev/shm 82
1.1.2.2.1 Ensure /dev/shm is tmpfs or a separate partition (Automated) 83
1.1.2.2.2 Ensure nodev option set on /dev/shm partition (Automated) 85
1.1.2.2.3 Ensure nosuid option set on /dev/shm partition (Automated) 87
1.1.2.2.4 Ensure noexec option set on /dev/shm partition (Automated) 89
1.1.2.3 Configure /home 91
1.1.2.3.1 Ensure separate partition exists for /home (Automated) 92
1.1.2.3.2 Ensure nodev option set on /home partition (Automated) 95
1.1.2.3.3 Ensure nosuid option set on /home partition (Automated) 97
1.1.2.4 Configure /var 99
1.1.2.4.1 Ensure separate partition exists for /var (Automated) 100
1.1.2.4.2 Ensure nodev option set on /var partition (Automated) 103
1.1.2.4.3 Ensure nosuid option set on /var partition (Automated) 105
1.1.2.5 Configure /var/tmp 107
1.1.2.5.1 Ensure separate partition exists for /var/tmp (Automated) 108
1.1.2.5.2 Ensure nodev option set on /var/tmp partition (Automated) 110
1.1.2.5.3 Ensure nosuid option set on /var/tmp partition (Automated) 112
1.1.2.5.4 Ensure noexec option set on /var/tmp partition (Automated) 114
1.1.2.6 Configure /var/log 116
1.1.2.6.1 Ensure separate partition exists for /var/log (Automated) 117
1.1.2.6.2 Ensure nodev option set on /var/log partition (Automated) 120
1.1.2.6.3 Ensure nosuid option set on /var/log partition (Automated) 122
1.1.2.6.4 Ensure noexec option set on /var/log partition (Automated) 124
1.1.2.7 Configure /var/log/audit 126
1.1.2.7.1 Ensure separate partition exists for /var/log/audit (Automated) 127
1.1.2.7.2 Ensure nodev option set on /var/log/audit partition (Automated) 130
1.1.2.7.3 Ensure nosuid option set on /var/log/audit partition (Automated) 132
1.1.2.7.4 Ensure noexec option set on /var/log/audit partition (Automated) 134
1.2 Package Management 136
1.2.1 Configure Package Repositories 137
1.2.1.1 Ensure GPG keys are configured (Manual) 138
1.2.1.2 Ensure gpgcheck is configured (Automated) 141
1.2.1.3 Ensure repo_gpgcheck is configured (Automated) 143
1.2.1.4 Ensure package manager repositories are configured (Manual) 146
1.2.1.5 Ensure weak dependencies are configured (Automated) 148
1.2.2 Configure Package Updates 151
1.2.2.1 Ensure updates, patches, and additional security software are installed (Manual) 152
1.3 Mandatory Access Control 154
1.3.1 Configure SELinux 155
1.3.1.1 Ensure SELinux is installed (Automated) 157
1.3.1.2 Ensure SELinux is not disabled in bootloader configuration (Automated) 159

1.3.1.3 Ensure SELinux policy is configured (Automated) 162
1.3.1.4 Ensure the SELinux mode is not disabled (Automated) 164
1.3.1.5 Ensure the SELinux mode is enforcing (Automated) 167
1.3.1.6 Ensure no unconfined services exist (Manual) 170
1.3.1.7 Ensure the MCS Translation Service (mcstrans) is not installed (Automated) 173
1.3.1.8 Ensure SETroubleshoot is not installed (Automated) 175
1.4 Configure Bootloader 177
1.4.1 Ensure bootloader password is set (Automated) 179
1.4.2 Ensure access to bootloader config is configured (Automated) 181
1.4.3 Ensure authentication is required when booting into rescue mode (Automated) 184
1.4.4 Ensure authentication is required when booting into emergency mode (Automated) 186
1.4.5 Ensure systemd ability to spawn an interactive boot process is disabled (Automated) 188
1.5 Configure Additional Process Hardening 190
1.5.1 Ensure fs.protected_hardlinks is configured (Automated) 192
1.5.2 Ensure fs.protected_symlinks is configured (Automated) 196
1.5.3 Ensure fs.suid_dumpable is configured (Automated) 200
1.5.4 Ensure kernel.dmesg_restrict is configured (Automated) 204
1.5.5 Ensure kernel.kptr_restrict is configured (Automated) 208
1.5.6 Ensure kernel.randomize_va_space is configured (Automated) 212
1.5.7 Ensure kernel.yama.ptrace_scope is configured (Automated) 216
1.5.8 Ensure core file size is configured (Automated) 221
1.5.9 Ensure systemd-coredump ProcessSizeMax is configured (Automated) 223
1.5.10 Ensure systemd-coredump Storage is configured (Automated) 226
1.6 Configure system wide crypto policy 229
1.6.1 Ensure system wide crypto policy is not legacy (Automated) 232
1.6.2 Ensure system wide crypto policy disables sha1 hash and signature support (Automated) 235
1.6.3 Ensure system wide crypto policy macs are configured (Automated) 238
1.6.4 Ensure system wide crypto policy disables cbc for ssh (Automated) 240
1.6.5 Ensure system wide crypto policy disables chacha20-poly1305 for ssh (Manual) 242
1.6.6 Ensure system wide crypto policy disables EtM for ssh (Manual) 245
1.6.7 Ensure system wide crypto policy enables post-quantum algorithms (Automated) 248
1.7 Configure Command Line Warning Banners 251
1.7.1 Ensure /etc/motd is configured (Automated) 252
1.7.2 Ensure /etc/issue is configured (Automated) 254
1.7.3 Ensure /etc/issue.net is configured (Automated) 256
1.7.4 Ensure access to /etc/motd is configured (Automated) 258
1.7.5 Ensure access to /etc/issue is configured (Automated) 260
1.7.6 Ensure access to /etc/issue.net is configured (Automated) 262
1.8 Configure GNOME Display Manager 264
1.8.1 Ensure GDM login banner is configured (Automated) 265
1.8.2 Ensure GDM disable-user-list is configured (Automated) 270
1.8.3 Ensure GDM screen lock is configured (Automated) 274
1.8.4 Ensure GDM automount is configured (Automated) 280
1.8.5 Ensure GDM autorun-never is configured (Automated) 285
1.8.6 Ensure XDMCP is not enabled (Automated) 288
1.8.7 Ensure Wayland is configured (Automated) 291

2 Services 293
2.1 Configure Server Services 294
2.1.1 Ensure autofs services are not in use (Automated) 295
2.1.2 Ensure avahi daemon services are not in use (Automated) 298
2.1.3 Ensure cockpit web services are not in use (Automated) 301
2.1.4 Ensure dhcp server services are not in use (Automated) 304
2.1.5 Ensure dns server services are not in use (Automated) 307
2.1.6 Ensure dnsmasq services are not in use (Automated) 310

2.1.7 Ensure samba file server services are not in use (Automated) 313
2.1.8 Ensure ftp server services are not in use (Automated) 316
2.1.9 Ensure message access server services are not in use (Automated) 319
2.1.10 Ensure network file system services are not in use (Automated) 322
2.1.11 Ensure print server services are not in use (Automated) 325
2.1.12 Ensure rpcbind services are not in use (Automated) 328
2.1.13 Ensure rsync services are not in use (Automated) 331
2.1.14 Ensure snmp services are not in use (Automated) 334
2.1.15 Ensure telnet server services are not in use (Automated) 337
2.1.16 Ensure tftp server services are not in use (Automated) 340
2.1.17 Ensure web proxy server services are not in use (Automated) 343
2.1.18 Ensure web server services are not in use (Automated) 346
2.1.19 Ensure mail transfer agents are configured for local-only mode (Automated) 350
2.1.20 Ensure a graphical user interface is not in use (Automated) 353
2.1.21 Ensure X window server services are not in use (Automated) 356
2.1.22 Ensure Xwayland is not in use (Automated) 358
2.1.23 Ensure only approved services are listening on a network interface (Manual) 361
2.2 Configure Client Services 364
2.2.1 Ensure ftp client is not installed (Automated) 365
2.2.2 Ensure ldap client is not installed (Automated) 367
2.2.3 Ensure telnet client is not installed (Automated) 369
2.2.4 Ensure tftp client is not installed (Automated) 371
2.3 Configure Time Synchronization 373
2.3.1 Ensure time synchronization is in use (Automated) 374
2.3.2 Ensure chrony is configured (Automated) 376
2.3.3 Ensure chrony is not run as the root user (Automated) 378
2.4 Job Schedulers 380
2.4.1 Configure cron 381
2.4.1.1 Ensure cron daemon is enabled and active (Automated) 382
2.4.1.2 Ensure access to /etc/crontab is configured (Automated) 384
2.4.1.3 Ensure access to /etc/cron.hourly is configured (Automated) 386
2.4.1.4 Ensure access to /etc/cron.daily is configured (Automated) 388
2.4.1.5 Ensure access to /etc/cron.weekly is configured (Automated) 390
2.4.1.6 Ensure access to /etc/cron.monthly is configured (Automated) 392
2.4.1.7 Ensure access to /etc/cron.yearly is configured (Automated) 394
2.4.1.8 Ensure access to /etc/cron.d is configured (Automated) 396
2.4.1.9 Ensure access to crontab is configured (Automated) 398
2.4.2 Configure at 402
2.4.2.1 Ensure access to at is configured (Automated) 403
3 Network 407
3.1 Configure Network Devices 408
3.1.1 Ensure IPv6 status is identified (Manual) 409
3.1.2 Ensure wireless interfaces are not available (Automated) 412
3.1.3 Ensure bluetooth services are not in use (Automated) 415
3.2 Configure Network Kernel Modules 418
3.2.1 Ensure dccp kernel module is not available (Automated) 419
3.2.2 Ensure tipc kernel module is not available (Automated) 422
3.2.3 Ensure rds kernel module is not available (Automated) 426
3.2.4 Ensure sctp kernel module is not available (Automated) 429
3.2.5 Ensure atm kernel module is not available (Automated) 433
3.2.6 Ensure can kernel module is not available (Automated) 437
3.3 Configure Network Kernel Parameters 441
3.3.1 Configure IPv4 parameters 442
3.3.1.1 Ensure net.ipv4.ip_forward is configured (Automated) 443
3.3.1.2 Ensure net.ipv4.conf.all.forwarding is configured (Automated) 447

3.3.1.3 Ensure net.ipv4.conf.default.forwarding is configured (Automated) 451
3.3.1.4 Ensure net.ipv4.conf.all.send_redirects is configured (Automated) 455
3.3.1.5 Ensure net.ipv4.conf.default.send_redirects is configured (Automated) 459
3.3.1.6 Ensure net.ipv4.icmp_ignore_bogus_error_responses is configured (Automated) ... 463
3.3.1.7 Ensure net.ipv4.icmp_echo_ignore_broadcasts is configured (Automated) 467
3.3.1.8 Ensure net.ipv4.conf.all.accept_redirects is configured (Automated) 471
3.3.1.9 Ensure net.ipv4.conf.default.accept_redirects is configured (Automated) 475
3.3.1.10 Ensure net.ipv4.conf.all.secure_redirects is configured (Automated) 479
3.3.1.11 Ensure net.ipv4.conf.default.secure_redirects is configured (Automated) 483
3.3.1.12 Ensure net.ipv4.conf.all.rp_filter is configured (Automated) 487
3.3.1.13 Ensure net.ipv4.conf.default.rp_filter is configured (Automated) 491
3.3.1.14 Ensure net.ipv4.conf.all.accept_source_route is configured (Automated) 495
3.3.1.15 Ensure net.ipv4.conf.default.accept_source_route is configured (Automated) 499
3.3.1.16 Ensure net.ipv4.conf.all.log_martians is configured (Automated) 503
3.3.1.17 Ensure net.ipv4.conf.default.log_martians is configured (Automated) 507
3.3.1.18 Ensure net.ipv4.tcp_syncookies is configured (Automated) 511
3.3.1.19 Ensure net.ipv4.conf.all.route_localnet is configured (Automated) 515
3.3.1.20 Ensure net.ipv4.conf.default.route_localnet is configured (Automated) 519
3.3.2 Configure IPv6 parameters 522
3.3.2.1 Ensure net.ipv6.conf.all.forwarding is configured (Automated) 523
3.3.2.2 Ensure net.ipv6.conf.default.forwarding is configured (Automated) 527
3.3.2.3 Ensure net.ipv6.conf.all.accept_redirects is configured (Automated) 531
3.3.2.4 Ensure net.ipv6.conf.default.accept_redirects is configured (Automated) 535
3.3.2.5 Ensure net.ipv6.conf.all.accept_source_route is configured (Automated) 539
3.3.2.6 Ensure net.ipv6.conf.default.accept_source_route is configured (Automated) 543
3.3.2.7 Ensure net.ipv6.conf.all.accept_ra is configured (Automated) 547
3.3.2.8 Ensure net.ipv6.conf.default.accept_ra is configured (Automated) 551

4 Host Based Firewall 555
4.1 Configure firewalld 556
4.1.1 Ensure firewalld is installed (Automated) 558
4.1.2 Ensure firewalld backend is configured (Automated) 560
4.1.3 Ensure firewalld.service is configured (Automated) 562
4.1.4 Ensure firewalld active zone target is configured (Automated) 564
4.1.5 Ensure firewalld loopback traffic is configured (Manual) 567
4.1.6 Ensure firewalld loopback source address traffic is configured (Manual) 569
4.1.7 Ensure firewalld services and ports are configured (Manual) 572
4.1.8 Ensure firewalld default zone is set to drop (Manual) 574
4.1.9 Ensure conflicting firewall utilities are not in use with firewalld (Automated) 576
4.1.10 Ensure firewalld logs denied packets (Automated) 578
4.1.11 Ensure firewalld blocks ICMP timestamp requests and replies (Automated) 579

5 Access Control 582
5.1 Configure SSH Server 583
5.1.1 Ensure access to /etc/ssh/sshd_config is configured (Automated) 586
5.1.2 Ensure access to /etc/ssh/sshd_config.d/ directory is configured (Automated) 588
5.1.3 Ensure access to SSH private host key files is configured (Automated) 590
5.1.4 Ensure access to SSH public host key files is configured (Automated) 594
5.1.5 Ensure sshd Ciphers are configured (Automated) 597
5.1.6 Ensure sshd KexAlgorithms is configured (Automated) 600
5.1.7 Ensure sshd MACs are configured (Automated) 603
5.1.8 Ensure sshd access is configured (Automated) 606
5.1.9 Ensure sshd Banner is configured (Automated) 609
5.1.10 Ensure sshd ClientAliveInterval and ClientAliveCountMax are configured (Automated) 611
5.1.11 Ensure sshd DisableForwarding is enabled (Automated) 614
5.1.12 Ensure sshd GSSAPIAuthentication is disabled (Automated) 616

5.1.13 Ensure sshd HostbasedAuthentication is disabled (Automated) 618
5.1.14 Ensure sshd IgnoreRhosts is enabled (Automated) 620
5.1.15 Ensure sshd LoginGraceTime is configured (Automated) 622
5.1.16 Ensure sshd LogLevel is configured (Automated) 624
5.1.17 Ensure sshd MaxAuthTries is configured (Automated) 626
5.1.18 Ensure sshd MaxStartups is configured (Automated) 628
5.1.19 Ensure sshd MaxSessions is configured (Automated) 630
5.1.20 Ensure sshd PermitEmptyPasswords is disabled (Automated) 632
5.1.21 Ensure sshd PermitRootLogin is disabled (Automated) 634
5.1.22 Ensure sshd PermitUserEnvironment is disabled (Automated) 636
5.1.23 Ensure sshd UsePAM is enabled (Automated) 638
5.2 Configure privilege escalation 640
5.2.1 Ensure sudo is installed (Automated) 641
5.2.2 Ensure access on sudoers configuration files is configured (Automated) 643
5.2.3 Ensure access on the /etc/sudoers.d directory is configured (Automated) 645
5.2.4 Ensure sudo commands use pty (Automated) 647
5.2.5 Ensure sudo log file exists (Automated) 650
5.2.6 Ensure users must provide password for escalation (Automated) 653
5.2.7 Ensure re-authentication for privilege escalation is not disabled globally (Automated) 655
5.2.8 Ensure sudo pwfeedback is configured (Automated) 657
5.2.9 Ensure sudo visiblepw is configured (Automated) 659
5.2.10 Ensure sudo timestamp_timeout is configured (Automated) 661
5.2.11 Ensure sudo timestamp_type is configured (Automated) 663
5.2.12 Ensure access to the su command is restricted (Automated) 665
5.3 Pluggable Authentication Modules 668
5.3.1 Configure PAM software packages 669
5.3.1.1 Ensure latest version of pam is installed (Automated) 670
5.3.1.2 Ensure latest version of authselect is installed (Automated) 672
5.3.1.3 Ensure latest version of libpwquality is installed (Automated) 674
5.3.2 Configure authselect 676
5.3.2.1 Ensure PAM is managed by an active authselect profile (Automated) 678
5.3.2.2 Ensure pam_faillock module is enabled (Automated) 681
5.3.2.3 Ensure pam_pwquality module is enabled (Automated) 684
5.3.2.4 Ensure pam_pwhistory module is enabled (Automated) 687
5.3.2.5 Ensure pam_unix module is enabled (Automated) 690
5.3.3 Configure PAM Arguments 693
5.3.3.1 Configure pam_faillock module 694
5.3.3.1.1 Ensure password failed attempts lockout is configured (Automated) 695
5.3.3.1.2 Ensure password unlock time is configured (Automated) 697
5.3.3.1.3 Ensure password failed attempts lockout includes root account (Automated) 700
5.3.3.2 Configure pam_pwquality module 704
5.3.3.2.1 Ensure password number of changed characters is configured (Automated) 705
5.3.3.2.2 Ensure password length is configured (Automated) 708
5.3.3.2.3 Ensure password complexity is configured (Manual) 711
5.3.3.2.4 Ensure password same consecutive characters is configured (Automated) 715
5.3.3.2.5 Ensure password maximum sequential characters is configured (Automated) 718
5.3.3.2.6 Ensure password dictionary check is enabled (Automated) 721
5.3.3.2.7 Ensure password quality is enforced for the root user (Automated) 723
5.3.3.3 Configure pam_pwhistory module 725
5.3.3.3.1 Ensure password history remember is configured (Automated) 727
5.3.3.3.2 Ensure password history is enforced for the root user (Automated) 729
5.3.3.3.3 Ensure pam_pwhistory includes use_authtok (Automated) 731
5.3.3.4 Configure pam_unix module 733
5.3.3.4.1 Ensure pam_unix does not include nullok (Automated) 734
5.3.3.4.2 Ensure pam_unix does not include remember (Automated) 736
5.3.3.4.3 Ensure pam_unix includes a strong password hashing algorithm (Automated) 738

5.3.3.4.4 Ensure pam_unix includes use_authtok (Automated) 741
5.4 User Accounts and Environment 743
5.4.1 Configure shadow password suite parameters 744
5.4.1.1 Ensure password expiration is configured (Automated) 745
5.4.1.2 Ensure minimum password days is configured (Manual) 748
5.4.1.3 Ensure password expiration warning days is configured (Automated) 751
5.4.1.4 Ensure strong password hashing algorithm is configured (Automated) 753
5.4.1.5 Ensure inactive password lock is configured (Automated) 756
5.4.1.6 Ensure all users last password change date is in the past (Automated) 758
5.4.2 Configure root and system accounts and environment 760
5.4.2.1 Ensure root is the only UID 0 account (Automated) 761
5.4.2.2 Ensure root is the only GID 0 account (Automated) 763
5.4.2.3 Ensure group root is the only GID 0 group (Automated) 765
5.4.2.4 Ensure root account access is controlled (Automated) 767
5.4.2.5 Ensure root path integrity (Automated) 769
5.4.2.6 Ensure root user umask is configured (Automated) 771
5.4.2.7 Ensure system accounts do not have a valid login shell (Automated) 773
5.4.2.8 Ensure accounts without a valid login shell are locked (Automated) 775
5.4.3 Configure user default environment 777
5.4.3.1 Ensure nologin is not listed in /etc/shells (Automated) 778
5.4.3.2 Ensure default user shell timeout is configured (Automated) 780
5.4.3.3 Ensure default user umask is configured in /etc/login.defs (Automated) 783
5.4.3.4 Ensure default user umask is configured (Automated) 785
6 Logging and Auditing 789
6.1 System Logging 790
6.1.1 Configure journald 791
6.1.1.1 Ensure journald service is active (Automated) 792
6.1.1.2 Ensure systemd-journal-remote is not in use (Automated) 794
6.1.1.3 Ensure journald is configured to send logs to rsyslog (Automated) 796
6.1.1.4 Ensure journald log file rotation is configured (Manual) 799
6.1.1.5 Ensure journald log file access is configured (Manual) 801
6.1.2 Configure rsyslog 803
6.1.2.1 Ensure rsyslog is installed (Automated) 805
6.1.2.2 Ensure rsyslog service is enabled and active (Automated) 807
6.1.2.3 Ensure rsyslog log file creation mode is configured (Automated) 809
6.1.2.4 Ensure access to rsyslog configuration files is configured (Automated) 813
6.1.2.5 Ensure access to the /etc/rsyslog.d directory is configured (Automated) 816
6.1.2.6 Ensure access to the rsyslog work directory is configured (Automated) 818
6.1.2.7 Ensure rsyslog is not configured as a log server (Automated) 821
6.1.2.8 Ensure rsyslog logging rules are configured (Manual) 825
6.1.2.9 Ensure rsyslog is configured to send logs to a remote log host (Manual) 828
6.1.2.10 Ensure a rsyslog TLS network stream driver is installed (Automated) 831
6.1.2.11 Ensure rsyslog network forwarding uses TLS (Automated) 833
6.1.2.12 Ensure rsyslog TLS CA certificate is configured (Manual) 837
6.1.2.13 Ensure rsyslog TLS client certificate is configured (Manual) 840
6.1.2.14 Ensure logrotate is configured (Manual) 843
6.1.3 Configure Logfiles 846
6.1.3.1 Ensure access to all logfiles has been configured (Automated) 847
6.2 System Auditing 852
6.2.1 Configure auditd Service 854
6.2.1.1 Ensure auditd packages are installed (Automated) 855
6.2.1.2 Ensure auditing for processes that start prior to auditd is enabled (Automated) 857
6.2.1.3 Ensure audit_backlog_limit is configured (Automated) 859
6.2.1.4 Ensure auditd service is enabled and active (Automated) 861
6.2.2 Configure Data Retention 863

6.2.2.1 Ensure audit log storage size is configured (Automated) 864
6.2.2.2 Ensure audit logs are not automatically deleted (Automated) 866
6.2.2.3 Ensure system is disabled when audit logs are full (Automated) 868
6.2.2.4 Ensure system warns when audit logs are low on space (Automated) 871
6.2.3 Configure auditd Rules 874
General guidance for auditd rules (applies to all recommendations in this section) 874
6.2.3.1 Ensure modification of the /etc/sudoers file is collected (Automated) 878
6.2.3.2 Ensure actions as another user are always logged (Automated) 881
6.2.3.3 Ensure events that modify the sudo log file are collected (Automated) 884
6.2.3.4 Ensure events that modify date and time information are collected (Automated) 888
6.2.3.5 Ensure events that modify sethostname and setdomainname are collected (Automated) 892
6.2.3.6 Ensure events that modify /etc/issue and /etc/issue.net are collected (Automated) .. 895
6.2.3.7 Ensure events that modify /etc/hosts and /etc/hostname are collected (Automated) 898
6.2.3.8 Ensure events that modify /etc/sysconfig/network and /etc/NetworkManager/system-connections/ are collected (Automated) 901
6.2.3.9 Ensure events that modify /etc/NetworkManager directory are collected (Automated) 905
6.2.3.10 Ensure use of privileged commands are collected (Automated) 907
6.2.3.11 Ensure unsuccessful file access attempts are collected (Automated) 912
6.2.3.12 Ensure events that modify /etc/group information are collected (Automated) 916
6.2.3.13 Ensure events that modify /etc/passwd information are collected (Automated) 918
6.2.3.14 Ensure events that modify /etc/shadow and /etc/gshadow are collected (Automated) 920
6.2.3.15 Ensure events that modify /etc/security/opasswd are collected (Automated) 923
6.2.3.16 Ensure events that modify /etc/nsswitch.conf file are collected (Automated) 926
6.2.3.17 Ensure events that modify /etc/pam.conf and /etc/pam.d/ information are collected (Automated) 928
6.2.3.18 Ensure discretionary access control permission modification events chmod,fchmod,fchmodat are collected (Automated) 931
6.2.3.19 Ensure discretionary access control permission modification events chown,fchown,lchown,fchownat are collected (Automated) 935
6.2.3.20 Ensure discretionary access control permission modification events setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr collected (Automated) 939
6.2.3.21 Ensure successful file system mounts are collected (Automated) 943
6.2.3.22 Ensure session initiation information is collected (Automated) 947
6.2.3.23 Ensure login and logout events are collected (Automated) 951
6.2.3.24 Ensure unlink file deletion events by users are collected (Automated) 954
6.2.3.25 Ensure rename file deletion events by users are collected (Automated) 958
6.2.3.26 Ensure events that modify the system's Mandatory Access Controls are collected (Automated) 962
6.2.3.27 Ensure successful and unsuccessful attempts to use the chcon command are collected (Automated) 966
6.2.3.28 Ensure successful and unsuccessful attempts to use the setfacl command are collected (Automated) 969
6.2.3.29 Ensure successful and unsuccessful attempts to use the chacl command are collected (Automated) 972
6.2.3.30 Ensure successful and unsuccessful attempts to use the usermod command are collected (Automated) 975
6.2.3.31 Ensure kernel module loading unloading and modification is collected (Automated) 978
6.2.3.32 Ensure kernel "init_module" and "finit_module" loading unloading and modification is collected (Automated) 981
6.2.3.33 Ensure kernel "delete_module" loading unloading and modification is collected (Automated) 985
6.2.3.34 Ensure the audit configuration is loaded regardless of errors (Automated) 989

6.2.3.35 Ensure the audit configuration is immutable (Automated) ... 991
6.2.3.36 Ensure the running and on disk configuration is the same (Automated) ... 994
6.2.4 Configure auditd File Access ... 996
6.2.4.1 Ensure the audit log file directory mode is configured (Automated) ... 997
6.2.4.2 Ensure audit log files mode is configured (Automated) ... 999
6.2.4.3 Ensure audit log files owner is configured (Automated) ... 1001
6.2.4.4 Ensure audit log files group owner is configured (Automated) ... 1003
6.2.4.5 Ensure audit configuration files mode is configured (Automated) ... 1005
6.2.4.6 Ensure audit configuration files owner is configured (Automated) ... 1007
6.2.4.7 Ensure audit configuration files group owner is configured (Automated) ... 1009
6.2.4.8 Ensure audit tools mode is configured (Automated) ... 1011
6.2.4.9 Ensure audit tools owner is configured (Automated) ... 1013
6.2.4.10 Ensure audit tools group owner is configured (Automated) ... 1015
6.3 Configure Integrity Checking ... 1017
6.3.1 Ensure AIDE is installed (Automated) ... 1018
6.3.2 Ensure filesystem integrity is regularly checked (Automated) ... 1020
6.3.3 Ensure cryptographic mechanisms are used to protect the integrity of audit tools (Automated) ... 1024
6.3.4 Ensure report_url is configured (Automated) ... 1028
7 System Maintenance ... 1031
7.1 Configure system file and directory access ... 1032
7.1.1 Ensure access to /etc/passwd is configured (Automated) ... 1033
7.1.2 Ensure access to /etc/passwd- is configured (Automated) ... 1035
7.1.3 Ensure access to /etc/group is configured (Automated) ... 1037
7.1.4 Ensure access to /etc/group- is configured (Automated) ... 1039
7.1.5 Ensure access to /etc/shadow is configured (Automated) ... 1041
7.1.6 Ensure access to /etc/shadow- is configured (Automated) ... 1043
7.1.7 Ensure access to /etc/gshadow is configured (Automated) ... 1045
7.1.8 Ensure access to /etc/gshadow- is configured (Automated) ... 1047
7.1.9 Ensure access to /etc/shells is configured (Automated) ... 1049
7.1.10 Ensure access to /etc/security/opasswd is configured (Automated) ... 1051
7.1.11 Ensure world writable files and directories are secured (Automated) ... 1053
7.1.12 Ensure no files or directories without an owner and a group exist (Automated) ... 1057
7.1.13 Ensure SUID and SGID files are reviewed (Manual) ... 1059
7.2 Local User and Group Settings ... 1061
7.2.1 Ensure accounts in /etc/passwd use shadowed passwords (Automated) ... 1062
7.2.2 Ensure /etc/shadow password fields are not empty (Automated) ... 1064
7.2.3 Ensure all groups in /etc/passwd exist in /etc/group (Automated) ... 1066
7.2.4 Ensure no duplicate UIDs exist (Automated) ... 1068
7.2.5 Ensure no duplicate GIDs exist (Automated) ... 1070
7.2.6 Ensure no duplicate user names exist (Automated) ... 1072
7.2.7 Ensure no duplicate group names exist (Automated) ... 1074
7.2.8 Ensure local interactive user home directories are configured (Automated) ... 1076
7.2.9 Ensure local interactive user dot files access is configured (Automated) ... 1079
Appendix: Summary Table ... 1082
Appendix: CIS Controls v7 IG 1 Mapped Recommendations ... 1107
Appendix: CIS Controls v7 IG 2 Mapped Recommendations ... 1113
Appendix: CIS Controls v7 IG 3 Mapped Recommendations ... 1125
Appendix: CIS Controls v7 Unmapped Recommendations ... 1137
Appendix: CIS Controls v8 IG 1 Mapped Recommendations ... 1138

Appendix: CIS Controls v8 IG 2 Mapped Recommendations 1145
Appendix: CIS Controls v8 IG 3 Mapped Recommendations 1157
Appendix: CIS Controls v8 Unmapped Recommendations 1169
Appendix: Change History 1170

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.

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.

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.

Target Technology Details

This document provides prescriptive guidance for establishing a secure configuration posture for AlmaLinux OS 9 systems running on x86_64 platforms.

This guide was developed and tested against AlmaLinux OS 9.8

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

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

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

Intended Audience

This benchmark is intended for system and application administrators, security specialists, auditors, help desk, and platform deployment personnel who plan to develop, deploy, assess, or secure solutions that incorporate Red Hat Enterprise Linux 9 on x86_64 platforms.

Consensus Guidance

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

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

Typographical Conventions

The following typographical conventions are used throughout this guide:

Convention	Meaning
`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).

Recommendation Definitions

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

Title

Concise description for the recommendation's intended configuration.

Assessment Status

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

Automated

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

Manual

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

Profile

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

Description

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

Rationale Statement

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

Impact Statement

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

Audit Procedure

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

Remediation Procedure

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

Default Value

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

References

Additional documentation relative to the recommendation.

CIS Critical Security Controls® (CIS Controls®)

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

Additional Information

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

Profile Definitions

The following configuration profiles are defined by this Benchmark:

- **Level 1 - Server**

  Items in this profile intend to:

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

  This profile is intended for servers.

- **Level 2 - Server**

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

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

  This profile is intended for servers.

- **Level 1 - Workstation**

  Items in this profile intend to:

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

  This profile is intended for workstations.

- **Level 2 - Workstation**

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

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

  This profile is intended for workstations.

Acknowledgements

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

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

Contributor
Ron Colvin
Dominic Pace
Mark Birch
Thomas Sjögren
James Trigg
Matthew Burket
Marcus Burghardt
Graham Eames
Ryan Jaynes
Simon John
Steve Cobrin

Editor
Jonathan Lewis Christopherson
Eric Pinnell
Gokhan Lus
Randie Bejar
Trevor Wood

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.

1.1.1 Configure Filesystem Kernel Modules

This section covers ensuring unneeded and legacy filesystem kernel modules are not available.

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`.

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.

Impact:

Removing support for the `cramfs` filesystem prevents mounting `cramfs`-formatted images. `cramfs` is used mainly in some embedded and initramfs contexts, so the impact on a typical server is minimal.

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:

```
# 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-
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
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.		●	●

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.

Impact:

Removing support for the `freevxfs` filesystem prevents mounting Veritas (VxFS) volumes. This filesystem is rarely used on Linux, so the impact on a typical server is minimal.

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:

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

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.		●	●

1.1.1.3 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.

Impact:

Removing support for the `hfs` filesystem prevents mounting legacy Apple HFS-formatted media.

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:

```
# 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

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.		●	●

1.1.1.4 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.

Impact:

Removing support for the `hfsplus` filesystem prevents mounting Apple HFS+-formatted media, such as Mac-formatted USB drives and disks.

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:

```
# 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

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.		●	●

1.1.1.5 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.

Impact:

Removing support for the `jffs2` filesystem prevents mounting JFFS2 (flash) volumes. This filesystem is used mainly on embedded devices, so the impact on a typical server is minimal.

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:

```
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

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.		●	●

1.1.1.6 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.

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

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.		●	●

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:

WARNING: `squashfs` is used by Snap packages and by some live/installer media and container images. Disabling it may break `snap` applications and those workflows. Confirm none are in use before disabling.

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:

```
#!/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:

```
# 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.		●	●

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:

Removing support for the `udf` filesystem prevents mounting UDF-formatted optical media, such as DVDs, and some USB media.

WARNING Microsoft Azure CSP requires the usage of udf

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.

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:

```
# 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.		●	●

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.

Impact:

Removing support for the `firewire-core` module disables IEEE 1394 (FireWire) device support on the system.

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:

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
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.		●	●

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.

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.

1. 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.

Note: Some systems may include usb-storage 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 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

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.		●	●

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.

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:

```
#!/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`.

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:

- 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

```
#!/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" = "overlayfs" ] && l_mod_name="overlay"
         [ "$l_dl" != "y" ] && f_module_fix
      else
         printf '%s\n' " - kernel module: \"$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" ""
}
```

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 from resource exhaustion and enables the use of mount options applicable to the directory's intended use. Users' data can be stored on separate partitions with stricter mount options. A user partition is a filesystem established for use by users that does not contain software for system operations.

The recommendations in this section are easiest 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.

Persistent storage requirement: Several mount points in this section -- `/home`, `/var`, `/var/tmp`, `/var/log`, and `/var/log/audit` -- hold data that must survive a reboot, so each must be a separate mount on **persistent (non-volatile) storage**; a `tmpfs` or `ramfs` (RAM-backed) filesystem does not satisfy the recommendation. A bind mount is acceptable only when its source is on persistent storage: `findmnt` reports the backing filesystem, so a bind mount onto a `tmpfs` location shows `tmpfs` and fails the check. (`/tmp` and `/dev/shm` are intended to be volatile and are exempt from this requirement.)

Btrfs subvolumes: 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.

Mounts not defined in `/etc/fstab`: The remediations in this section show adding mount options to `/etc/fstab`. If a mount point is not configured there -- for example on systems using encrypted volumes (LUKS), where the entry appears as a `/dev/mapper/<name>` device, or where the mount is defined by a systemd `.mount` unit (as `/tmp` may be via `tmp.mount`) -- apply the required options wherever that mount is actually defined (the corresponding `/etc/fstab` line for the mapper device, or the `Options=` line of the systemd `.mount` unit), then run `mount -o remount <mount point>` (or reboot) to apply.

Repartitioning an already-installed system (this may require the system to be in single-user mode):

1. Mount the new partition to a temporary mountpoint, e.g. `mount /dev/sda2 /mnt`
2. Copy data from the original partition to the new partition, e.g. `cp -a /var/tmp/* /mnt`
3. Verify that all data is present on the new partition, e.g. `ls -la /mnt`
4. Unmount the new partition, e.g. `umount /mnt`
5. Remove the data from the original directory -- otherwise it will continue to consume space in the old partition, which is masked once the new filesystem is mounted, e.g. `rm -Rf /var/tmp/*`
6. Mount the new partition to the desired mountpoint, e.g. `mount /dev/sda2 /var/tmp`

7. Update `/etc/fstab` with the new mountpoint, e.g. `/dev/sda2 /var/tmp xfs defaults,rw,nosuid,nodev,noexec,relatime 0 0`

When repartitioning a mount whose service should be quiesced first -- for example `/var/log/audit` while `auditd` is running -- perform the migration from emergency mode (so the service is not writing to the directory), then return to multi-user mode.

1.1.2.1 Configure /tmp

This section covers configuring `/tmp` as a separate partition or `tmpfs` and applying the `nodev`, `nosuid`, and `noexec` mount options.

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.

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.

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
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.	●	●	●
v7	14.6 Protect Information through Access Control Lists Protect all information stored on systems with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.	●	●	●

1.1.2.1.2 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`.

Impact:

Setting the `nodev` option prevents device files from being interpreted on the `/tmp` filesystem, which is not expected to contain device files; this typically has no operational impact.

Audit:

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

Run the following command to verify the `nodev` option is set:

```
# 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

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.	●	●	●

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`.

Impact:

Setting the `nosuid` option prevents `setuid` and `setgid` bits from taking effect on the `/tmp` filesystem, which is not expected to contain privileged executables; this typically has no operational impact.

Audit:

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

Run the following command to verify the `nosuid` option is set:

```
# 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

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.	●	●	●

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 the `noexec` mount option is set.

Run the following command to verify the `noexec` option is set:

```
# 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
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.	●	●	●

1.1.2.2 Configure /dev/shm

This section covers configuring `/dev/shm` as a separate partition or `tmpfs` and applying the `nodev`, `nosuid`, and `noexec` mount options.

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

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`.

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.	●	●	●

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.

Impact:

Setting the `nodev` option prevents device files from being interpreted on the `/dev/shm` filesystem, which is not expected to contain device files; this typically has no operational impact.

Audit:

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

Run the following command to verify 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

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.	●	●	●

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.

Impact:

Setting the `nosuid` option prevents `setuid` and `setgid` bits from taking effect on the `/dev/shm` filesystem, which is not expected to contain privileged executables; this typically has no operational impact.

Audit:

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

Run the following command to verify 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

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.	●	●	●

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 the `noexec` mount option is set.

Run the following command to verify 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
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.	●	●	●

1.1.2.3 Configure /home

This section covers configuring `/home` as a separate partition and applying the `nodev` and `nosuid` mount options.

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
}
```

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.

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.

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.

`/home` stores persistent user data, so it must reside on persistent (non-volatile) storage.

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:

Verify that `/home` is a separate mount point backed by persistent (non-volatile) storage.

Run the following command:

```
# findmnt -knf -o TARGET,SOURCE,FSTYPE /home
```

`/home` should be returned as a separate mount whose `FSTYPE` is a persistent filesystem (for example `ext4`, `xfs`, or `btrfs`) -- not `tmpfs` or `ramfs`.

Example output:

```
/home /dev/sdb ext4
```

Note: The mount must be on persistent (non-volatile) storage; a bind mount is acceptable only if its source is persistent (see the section overview).

Remediation:

Configure `/home` as a separate mount on persistent (non-volatile) storage.

New installations: During installation, create a custom partitioning layout and assign a separate partition or logical volume to `/home`.

Existing systems: Create a new filesystem for `/home` and add it to the mount table. For example, using LVM (adjust the volume group, size, and device to suit the host):

```
# lvcreate -L <size> -n home <vg_name>
# mkfs.ext4 /dev/<vg_name>/home
```

Add the mount to `/etc/fstab`:

```
/dev/<vg_name>/home  /home  ext4  defaults,nodev,nosuid  0 0
```

Migrate the existing contents into the new filesystem and activate the mount using the repartitioning steps in the section overview, then run `mount /home` (or reboot).

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/

CIS Controls:

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

1.1.2.3.2 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`.

Impact:

Setting the `nodev` option prevents device files from being interpreted on the `/home` filesystem, which is not expected to contain device files; this typically has no operational impact.

Audit:

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

Run the following command to verify the `nodev` option is set:

```
# 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.

CIS Controls:

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

1.1.2.3.3 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`.

Impact:

Setting the `nosuid` option prevents `setuid` and `setgid` bits from taking effect on the `/home` filesystem, which is not expected to contain privileged executables; this typically has no operational impact.

Audit:

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

Run the following command to verify the `nosuid` option is set:

```
# 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.

CIS Controls:

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

1.1.2.4 Configure /var

This section covers configuring `/var` as a separate partition and applying the `nodev` and `nosuid` mount options.

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

1.1.2.4.1 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 `nodev` and `nosuid`. These options limit an attacker's ability to create exploits on the system. `noexec` is not applied to `/var`, which legitimately hosts executable content used by the package manager and services, so restricting execution there would break normal operation. Other options allow for specific behavior. See `man mount` for exact details regarding filesystem-independent and filesystem-specific options.

For example, a `setuid` program is normally patched by replacing the file -- the package manager writes a new file and renames it into place, creating a new inode -- not by modifying the running binary. An attacker who pre-created a hard link to that program (on the same filesystem) keeps a working copy of the old, vulnerable binary after it is patched: their hard link still references the original inode, so nothing is "broken." Because a hard link cannot span filesystems, placing `/var` on a separate partition helps prevent this by keeping world-writable locations off the filesystem that holds the system's `setuid` programs.

`/var` stores variable system data that must persist, so it must reside on persistent (non-volatile) storage.

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:

Verify that `/var` is a separate mount point backed by persistent (non-volatile) storage.

Run the following command:

```
# findmnt -knf -o TARGET,SOURCE,FSTYPE /var
```

`/var` should be returned as a separate mount whose `FSTYPE` is a persistent filesystem (for example `ext4`, `xfs`, or `btrfs`) -- not `tmpfs` or `ramfs`.

Example output:

```
/var /dev/sdb ext4
```

Note: The mount must be on persistent (non-volatile) storage; a bind mount is acceptable only if its source is persistent (see the section overview).

Remediation:

Configure `/var` as a separate mount on persistent (non-volatile) storage.

New installations: During installation, create a custom partitioning layout and assign a separate partition or logical volume to `/var`.

Existing systems: Create a new filesystem for `/var` and add it to the mount table. For example, using LVM (adjust the volume group, size, and device to suit the host):

```
# lvcreate -L <size> -n var <vg_name>
# mkfs.ext4 /dev/<vg_name>/var
```

Add the mount to `/etc/fstab`:

```
/dev/<vg_name>/var  /var  ext4  defaults,nodev,nosuid  0 0
```

Migrate the existing contents into the new filesystem and activate the mount using the repartitioning steps in the section overview, then run `mount /var` (or reboot).

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/

CIS Controls:

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

1.1.2.4.2 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`.

Impact:

Setting the `nodev` option prevents device files from being interpreted on the `/var` filesystem, which is not expected to contain device files; this typically has no operational impact.

Audit:

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

Run the following command to verify the `nodev` option is set:

```
# 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.

CIS Controls:

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

1.1.2.4.3 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`.

Impact:

Setting the `nosuid` option prevents `setuid` and `setgid` bits from taking effect on the `/var` filesystem, which is not expected to contain privileged executables; this typically has no operational impact.

Audit:

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

Run the following command to verify the `nosuid` option is set:

```
# 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.

CIS Controls:

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

1.1.2.5 Configure /var/tmp

This section covers configuring `/var/tmp` as a separate partition and applying the `nodev`, `nosuid`, and `noexec` mount options.

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

1.1.2.5.1 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.

Files in `/var/tmp` are expected to be preserved between reboots (FHS 3.0 §5.15), so it must reside on persistent (non-volatile) storage.

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:

Verify that `/var/tmp` is a separate mount point backed by persistent (non-volatile) storage.

Run the following command:

```
# findmnt -knf -o TARGET,SOURCE,FSTYPE /var/tmp
```

`/var/tmp` should be returned as a separate mount whose `FSTYPE` is a persistent filesystem (for example `ext4`, `xfs`, or `btrfs`) -- not `tmpfs` or `ramfs`.

Example output:

```
/var/tmp /dev/sdb ext4
```

Note: The mount must be on persistent (non-volatile) storage; a bind mount is acceptable only if its source is persistent (see the section overview).

Remediation:

Configure `/var/tmp` as a separate mount on persistent (non-volatile) storage.

New installations: During installation, create a custom partitioning layout and assign a separate partition or logical volume to `/var/tmp`.

Existing systems: Create a new filesystem for `/var/tmp` and add it to the mount table. For example, using LVM (adjust the volume group, size, and device to suit the host):

```
# lvcreate -L <size> -n var_tmp <vg_name>
# mkfs.ext4 /dev/<vg_name>/var_tmp
```

Add the mount to `/etc/fstab`:

```
/dev/<vg_name>/var_tmp  /var/tmp  ext4  defaults,nodev,nosuid,noexec  0 0
```

Migrate the existing contents into the new filesystem and activate the mount using the repartitioning steps in the section overview, then run `mount /var/tmp` (or reboot).

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

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`.

Impact:

Setting the `nodev` option prevents device files from being interpreted on the `/var/tmp` filesystem, which is not expected to contain device files; this typically has no operational impact.

Audit:

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

Run the following command to verify the `nodev` option is set:

```
# 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
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.	●	●	●

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`.

Impact:

Setting the `nosuid` option prevents `setuid` and `setgid` bits from taking effect on the `/var/tmp` filesystem, which is not expected to contain privileged executables; this typically has no operational impact.

Audit:

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

Run the following command to verify the `nosuid` option is set:

```
# 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
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.	●	●	●

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 the `noexec` mount option is set.

Run the following command to verify the `noexec` option is set:

```
# 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)
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.	●	●	●

1.1.2.6 Configure /var/log

This section covers configuring `/var/log` as a separate partition and applying the `nodev`, `nosuid`, and `noexec` mount options.

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

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.

`/var/log` stores system logs that must persist for operational and audit purposes, so it must reside on persistent (non-volatile) storage.

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:

Verify that `/var/log` is a separate mount point backed by persistent (non-volatile) storage.

Run the following command:

```
# findmnt -knf -o TARGET,SOURCE,FSTYPE /var/log
```

`/var/log` should be returned as a separate mount whose `FSTYPE` is a persistent filesystem (for example `ext4`, `xfs`, or `btrfs`) -- not `tmpfs` or `ramfs`.

Example output:

```
/var/log /dev/sdb ext4
```

Note: The mount must be on persistent (non-volatile) storage; a bind mount is acceptable only if its source is persistent (see the section overview).

Remediation:

Configure `/var/log` as a separate mount on persistent (non-volatile) storage.

New installations: During installation, create a custom partitioning layout and assign a separate partition or logical volume to `/var/log`.

Existing systems: Create a new filesystem for `/var/log` and add it to the mount table. For example, using LVM (adjust the volume group, size, and device to suit the host):

```
# lvcreate -L <size> -n var_log <vg_name>
# mkfs.ext4 /dev/<vg_name>/var_log
```

Add the mount to `/etc/fstab`:

```
/dev/<vg_name>/var_log  /var/log  ext4  defaults,nodev,nosuid,noexec  0 0
```

Migrate the existing contents into the new filesystem and activate the mount using the repartitioning steps in the section overview, then run `mount /var/log` (or reboot).

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/

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.		●	●

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`.

Impact:

Setting the `nodev` option prevents device files from being interpreted on the `/var/log` filesystem, which is not expected to contain device files; this typically has no operational impact.

Audit:

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

Run the following command to verify the `nodev` option is set:

```
# 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
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.	●	●	●

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`.

Impact:

Setting the `nosuid` option prevents `setuid` and `setgid` bits from taking effect on the `/var/log` filesystem, which is not expected to contain privileged executables; this typically has no operational impact.

Audit:

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

Run the following command to verify the `nosuid` option is set:

```
# 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
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.	●	●	●

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 the `noexec` mount option is set.

Run the following command to verify the `noexec` option is set:

```
# 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)
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.	●	●	●

1.1.2.7 Configure /var/log/audit

This section covers configuring `/var/log/audit` as a separate partition and applying the `nodev`, `nosuid`, and `noexec` mount options.

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

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.

`/var/log/audit` stores audit records that must persist for security and compliance, so it must reside on persistent (non-volatile) storage.

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:

Verify that `/var/log/audit` is a separate mount point backed by persistent (non-volatile) storage.

Run the following command:

```
# findmnt -knf -o TARGET,SOURCE,FSTYPE /var/log/audit
```

`/var/log/audit` should be returned as a separate mount whose `FSTYPE` is a persistent filesystem (for example `ext4`, `xfs`, or `btrfs`) -- not `tmpfs` or `ramfs`.

Example output:

```
/var/log/audit /dev/sdb ext4
```

Note: The mount must be on persistent (non-volatile) storage; a bind mount is acceptable only if its source is persistent (see the section overview).

Remediation:

Configure `/var/log/audit` as a separate mount on persistent (non-volatile) storage.

New installations: During installation, create a custom partitioning layout and assign a separate partition or logical volume to `/var/log/audit`.

Existing systems: Create a new filesystem for `/var/log/audit` and add it to the mount table. For example, using LVM (adjust the volume group, size, and device to suit the host):

```
# lvcreate -L <size> -n var_log_audit <vg_name>
# mkfs.ext4 /dev/<vg_name>/var_log_audit
```

Add the mount to `/etc/fstab`:

```
/dev/<vg_name>/var_log_audit  /var/log/audit  ext4
defaults,nodev,nosuid,noexec  0 0
```

Migrate the existing contents into the new filesystem and activate the mount using the repartitioning steps in the section overview, then run `mount /var/log/audit` (or reboot).

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/

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.		●	●

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`.

Impact:

Setting the `nodev` option prevents device files from being interpreted on the `/var/log/audit` filesystem, which is not expected to contain device files; this typically has no operational impact.

Audit:

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

Run the following command to verify the `nodev` option is set:

```
# 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
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.	●	●	●

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`.

Impact:

Setting the `nosuid` option prevents `setuid` and `setgid` bits from taking effect on the `/var/log/audit` filesystem, which is not expected to contain privileged executables; this typically has no operational impact.

Audit:

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

Run the following command to verify the `nosuid` option is set:

```
# 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)
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.	●	●	●

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 the `noexec` mount option is set.

Run the following command to verify the `noexec` option is set:

```
# 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
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 reveal the patched exploitable entry points to the public. Public knowledge of these exploits cans your organization more vulnerable to malicious actors attempting to gain entry to your system's data.

Software updates often offer new and improved features and speed enhancements

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

1.2.1 Configure Package Repositories

This section covers configuring the system's package repositories to use trusted sources with GPG signature verification.

Software is installed and updated from package repositories. Configuring those repositories correctly -- using trusted sources and verifying the authenticity and integrity of both the repository metadata and the packages themselves -- helps ensure a system receives legitimate software and updates and is not exposed to tampered or malicious packages.

The recommendations in this section cover the trust foundations for `dnf`/RPM package management: the GPG keys used to verify package signatures, package signature checking (`gpgcheck`), repository metadata signature checking (`repo_gpgcheck`), the set of configured repositories, and the handling of weak dependencies. Settings may be applied globally in `/etc/dnf/dnf.conf` and per-repository in files under the repository drop-in directories (for example `/etc/yum.repos.d`); per-repository settings take precedence over the global configuration.

1.2.1.1 Ensure GPG keys are configured (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

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

Rationale:

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

Audit:

List all GPG key URLs

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

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

List installed GPG keys

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

Example:

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

Example output:

```
RPM: gpg-pubkey-fd431d51-4ae0493b
Packager: Red Hat, Inc. (release key 2) <security@redhat.com>
Summary: Red Hat, Inc. (release key 2) <security@redhat.com> public key
Creation date: 2009-10-22
Key ID: fd431d51

RPM: gpg-pubkey-5a6340b3-6229229e
Packager: Red Hat, Inc. (auxiliary key 3) <security@redhat.com>
Summary: Red Hat, Inc. (auxiliary key 3) <security@redhat.com> public key
Creation date: 2022-03-09
Key ID: 5a6340b3
```

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

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

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

Query locally available GPG keys

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

To find where these keys come from run:

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

Remediation:

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

References:

1. GRID: NIX-00000076
2. NIST SP 800-53 Rev. 5: SI-2

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.	●	●	●

1.2.1.2 Ensure gpgcheck is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `gpgcheck` option controls verifying package signatures after download. This option is configurable as a global option in the main section of `/etc/dnf/dnf.conf` and a per repository option in individual files in the `/etc/yum.repos.d/*` directory.

The option is enabled if `gpgcheck` is set to `1`, `true`, or `yes`. The option is disabled if `gpgcheck` is set to `0`, `false`, or `no`. If an invalid option is set, e.g. `gpgcheck=2`, the global option will be used.

Settings in files in the `/etc/yum.repos.d/` directory take precedence over the global configuration.

Rationale:

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

Audit:

Verify that `gpgcheck` is enabled in the global DNF configuration and is not disabled by any repository.

Run the following command and verify that `gpgcheck` is set to `1` (the values `true` and `yes` are also accepted):

```
# dnf config-manager --dump | grep -Pi -- '^\h*gpgcheck\h*='
```

If `dnf config-manager` is not available, inspect the global configuration file directly:

```
# grep -Pi -- '^\h*gpgcheck\h*=' /etc/dnf/dnf.conf
```

A value of `1` should be returned.

Then verify that no repository disables `gpgcheck`. Run the following to list any repository (in any configured `reposdir` directory) that sets `gpgcheck=0`:

```
# grep -Pris -- '^\h*gpgcheck\h*=\h*(0|false|no)\b' /etc/yum.repos.d
/etc/yum/repos.d /etc/distro.repos.d 2>/dev/null
```

Nothing should be returned.

Remediation:

Set `gpgcheck=1` in the `[main]` section of `/etc/dnf/dnf.conf`:

```
[main]
gpgcheck=1
```

Ensure no repository overrides this with `gpgcheck=0`. In any repository definition file under `/etc/yum.repos.d` (or another configured `reposdir` directory) that sets `gpgcheck=0`, change it to `gpgcheck=1` (or remove the line so the repository inherits the global setting).

Default Value:

gpgcheck=1

References:

1. GRID: NIX-00000077
2. NIST SP 800-53 Revision 5 :: CM-14
3. STIG ID: RHEL-08-010370 | RULE ID: SV-230264r1017377 | CAT I
4. STIG ID: RHEL-09-214015 | RULE ID: SV-257820r1044878 | CAT I
5. STIG ID: RHEL-09-214025 | RULE ID: SV-257822r1044880 | CAT I

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.	●	●	●

1.2.1.3 Ensure repo_gpgcheck is configured (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

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

Rationale:

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

Impact:

Not all repositories support `repo_gpgcheck` (for example, some Red Hat repositories do not). If `repo_gpgcheck` is enabled for a repository whose repodata is not signed, package installation or updates from that repository will fail.

Setting `repo_gpgcheck=0` for a specific repository that does not support signed repodata is an acceptable exception and does not make the system non-compliant. Be aware, however, that the metadata for any repository with `repo_gpgcheck=0` is not signature-verified, so limit this to trusted repositories that genuinely do not support the option. Determining which configured repositories support `repo_gpgcheck` requires review by the administrator.

Audit:

Verify that `repo_gpgcheck` is enabled in the global DNF configuration.

Run the following command and verify that `repo_gpgcheck` is set to `1` (the values `true` and `yes` are also accepted):

```
# dnf config-manager --dump | grep -Pi -- '^\h*repo_gpgcheck\h*='
```

If `dnf config-manager` is not available, inspect the global configuration file directly:

```
# grep -Pi -- '^\h*repo_gpgcheck\h*=' /etc/dnf/dnf.conf
```

A value of `1` should be returned.

Per-repository exceptions: Repositories may override the global setting in their definition files under the repository drop-in directories (by default `/etc/yum.repos.d`, `/etc/yum/repos.d`, and `/etc/distro.repos.d`). A repository that sets `repo_gpgcheck=0` because it does not support signed repodata is acceptable; the metadata for such a repository is not signature-verified, so confirm each is trusted and genuinely unable to support the option. Run the following to list any repositories that disable `repo_gpgcheck`:

```
# grep -Pris -- '^\h*repo_gpgcheck\h*=\h*(0|false|no)\b' /etc/yum.repos.d
/etc/yum/repos.d /etc/distro.repos.d 2>/dev/null
```

Remediation:

Global configuration:

Edit `/etc/dnf/dnf.conf` and set the following in the `[main]` section:

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

Per-repository exceptions:

For any repository that does not support signed repodata, override the global setting in that repository's definition file under `/etc/yum.repos.d` (or another repository drop-in directory) by setting `repo_gpgcheck=0` for that repository. Per-repository settings take precedence over the global configuration. Limit exceptions to trusted repositories that genuinely do not support the option; their metadata will not be signature-verified.

References:

1. GRID: NIX-00000079
2. NIST SP 800-53 Rev. 5: SI-2

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.	●	●	●

1.2.1.4 Ensure package manager repositories are configured (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

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

Rationale:

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

Audit:

Verify that the configured package manager repositories point to trusted, authorized sources.

Run the following to list the enabled repositories:

```
# dnf repolist
```

Example output:

```
repo id                     repo name
fedora                      Fedora $releasever - $basearch
updates                     Fedora $releasever - $basearch - Updates
```

The repository id, name, and status shown by `dnf repolist` do not reveal the source URL. Inspect each repository definition and verify that its `baseurl` (or `mirrorlist`/`metalink`) resolves to a trusted source authorized by site policy -- an attacker who alters a `baseurl` can redirect the system to a malicious package source without changing the visible id or name:

```
# grep -E -- '^\s*(baseurl|mirrorlist|metalink)\s*=' /etc/yum.repos.d/*.repo
/etc/dnf/dnf.conf
```

Example output:

```
/etc/yum.repos.d/fedora.repo:metalink=https://mirrors.fedoraproject.org/metal
ink?repo=fedora-$releasever&arch=$basearch
/etc/yum.repos.d/fedora-
updates.repo:metalink=https://mirrors.fedoraproject.org/metalink?repo=updates
-released-f$releasever&arch=$basearch
/etc/yum.repos.d/local.repo:baseurl=https://repo.example.com/local/
```

Review any other configured `reposdir` directories as well. Confirm every configured source is expected and authorized for this system.

Note: Package and metadata integrity against a redirected or spoofed source is enforced by the `gpgcheck` and `repo_gpgcheck` recommendations -- with both enabled, a source serving tampered or unsigned content is rejected.

Remediation:

Configure your package manager repositories according to site policy.

References:

1. GRID: NIX-00000078
2. NIST SP 800-53 Rev. 5: SI-2

Additional Information:

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

CIS Controls:

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

1.2.1.5 Ensure weak dependencies are configured (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

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

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

Rationale:

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

Impact:

Weak dependencies (`Recommends`) provide optional functionality that is not required for core operation. With `install_weak_deps` disabled, those recommended packages are no longer installed automatically, so some optional features may be unavailable; any recommended package that is actually needed must be installed explicitly. The setting applies to future transactions only -- weak dependencies already installed are not removed.

Audit:

Run the following command and verify that `install_weak_deps` is disabled in the global configuration:

```
# grep -Pi -- '^\h*install_weak_deps\h*=\h*(0|false|no)\b' /etc/dnf/dnf.conf
```

One of the following should be returned: `install_weak_deps=0`, `install_weak_deps=false`, or `install_weak_deps=no`.

Example output:

```
install_weak_deps=0
```

Also verify that `install_weak_deps` is not set to an enabled value (`1`, `true`, or `yes`) elsewhere in the file:

```
# grep -Pi -- '^\h*install_weak_deps\h*=\h*(1|true|yes)\b' /etc/dnf/dnf.conf
```

Nothing should be returned.

Note: If no value is configured, the recommendation is not met -- the default value is enabled.

Remediation:

Edit `/etc/dnf/dnf.conf` and set `install_weak_deps=0` in the `[main]` section.

```
[main]
install_weak_deps=0
```

Example script:

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

{
   if grep -Pq '^install_weak_deps' /etc/dnf/dnf.conf; then
      sed -i 's/^install_weak_deps\s*=\s*.*/install_weak_deps=0/'
/etc/dnf/dnf.conf
   else
      printf '%s\n' "" "install_weak_deps=0" >> /etc/dnf/dnf.conf
   fi
}
```

Default Value:

```
install_weak_deps=True
```

References:

1. GRID: NIX-00000082
2. https://docs.fedoraproject.org/en-US/packaging-guidelines/WeakDependencies/

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.		●	●

1.2.2 Configure Package Updates

This section covers ensuring installed software is kept current with vendor updates and security patches.

Keeping installed software current with vendor updates and security patches is a core element of system maintenance. Patch-management procedures vary between enterprises: large environments may operate a local update server or mirror in place of the distribution's servers, while a single system may update directly. Updates may be applied automatically or manually depending on the site's patch-management policy, and organizations often test patches on non-production systems before rolling them out to production.

Outdated software is exposed to publicly known vulnerabilities; the release notes accompanying an update frequently disclose the very weaknesses it fixes, which attackers can then target on unpatched systems. Applying updates promptly, per a defined policy, reduces this exposure.

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

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

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Security vulnerabilities and functional improvements are regularly addressed through software updates and patches.

Rationale:

Unpatched software can expose systems to known vulnerabilities that may be exploited by attackers. It is recommended that operating system updates be performed on enterprise assets through automated patch management on a monthly, or more frequent, basis.

Impact:

Applying updates may require restarting affected services or rebooting the system (for kernel and core-library updates) to take effect, causing planned downtime. Updates can also change behavior or introduce regressions, which is why testing on non-production systems and following a defined patch-management policy is recommended.

Audit:

Run the following command and verify there are no updates or patches to install:

```
# dnf check-update
```

Check to make sure no system reboot is required:

```
# dnf needs-restarting -r
```

Remediation:

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

The following command will install all available updates:

```
# dnf update
```

Once the update process is complete, verify if reboot is required to load changes.

```
# dnf needs-restarting -r
```

References:

1. GRID: NIX-00000081
2. NIST SP 800-53 Rev. 5: SI-2

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.

1.3.1 Configure SELinux

This section covers installing, enabling, and enforcing SELinux, and removing components that weaken its policy.

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

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

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

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

Two such policies have been designed for use with Fedora Family Linux distributions and are included with the system:

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

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

Note:

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

References:

- NSA SELinux resources:
  - https://www.nsa.gov/Research/Technical-Papers-Brochures/smdsearch14229/selinux
- Fedora SELinux resources:
  - Getting started with SELinux: https://docs.fedoraproject.org/en-US/quick-docs/getting-started-with-selinux
  - User Guide: https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html-single/using_selinux/index
- SELinux Project web page and wiki:
  - http://www.selinuxproject.org

1.3.1.1 Ensure SELinux is installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

SELinux provides Mandatory Access Control.

Rationale:

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

Audit:

Verify SELinux is installed.

Run the following command:

```
# rpm -q libselinux

libselinux-<version>
```

Remediation:

Run the following command to install `SELinux`:

```
# dnf install libselinux
```

References:

1. GRID: NIX-00000090
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.	●	●	●

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

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

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

Rationale:

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

Impact:

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

Audit:

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

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

Nothing should be returned

Run the following command to verify the `selinux=0` and `enforcing=0` parameters were not configured by the deprecated `grub2-mkconfig` command:

```
# grep -Prs --
'^\h*([^#\n\r]+\h+)?kernelopts=([^#\n\r]+\h+)?(selinux|enforcing)=0\b'
/boot/grub2 /boot/efi
```

Nothing should be returned

Remediation:

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

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

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

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

References:

1. GRID: NIX-00000091
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

Additional Information:

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

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

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

CIS Controls:

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

1.3.1.3 Ensure SELinux policy is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

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

Rationale:

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

Audit:

Run the following command to verify SELinux policy is set to `targeted` or `mls`:

```
# grep -Psi -- '^\h*SELINUXTYPE\h*=\h*(targeted|mls)\b' /etc/selinux/config
```

Example output:

```
SELINUXTYPE=targeted
```

Remediation:

Edit `/etc/selinux/config` and set the SELINUXTYPE line to `targeted` or `mls`:

Example:

```
SELINUXTYPE=targeted
```

References:

1. GRID: NIX-00000092
2. NIST SP 800-53 Rev. 5: AC-3, MP-2

Additional Information:

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

CIS Controls:

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

1.3.1.4 Ensure the SELinux mode is not disabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

SELinux can run in one of three modes: disabled, permissive, or enforcing:

- `Enforcing` - Is the default, and recommended, mode of operation; in enforcing mode SELinux operates normally, enforcing the loaded security policy on the entire system.
- `Permissive` - The system acts as if SELinux is enforcing the loaded security policy, including labeling objects and emitting access denial entries in the logs, but it does not actually deny any operations. While not recommended for production systems, permissive mode can be helpful for SELinux policy development.
- `Disabled` - disables most of the SELinux kernel and application code, leaving the system running without any SELinux protection.

Note: You can set individual domains to permissive mode while the system runs in enforcing mode. For example, to make the httpd_t domain permissive:

```
# semanage permissive -a httpd_t
```

Rationale:

Running SELinux in disabled mode is strongly discouraged; not only does the system avoid enforcing the SELinux policy, it also avoids labeling any persistent objects such as files, making it difficult to enable SELinux in the future.

Impact:

Enabling SELinux from a `disabled` state (to either `enforcing` or `permissive`) triggers a full filesystem relabel on the next reboot, which can take a significant amount of time on large filesystems, and requires a reboot to take effect. Plan the change during a maintenance window.

Audit:

1. Run the following command to verify SELinux's configured mode:

```
# grep -Pi -- '^\h*SELINUX=(enforcing|permissive)\b' /etc/selinux/config
```

Verify the output is:

```
SELINUX=enforcing
-OR-
SELINUX=permissive
```

2. Run the following command to verify SELinux's current running mode:

```
# getenforce
```

Verify the output is:

```
Enforcing
-OR-
Permissive
```

Remediation:

1. Edit `/etc/selinux/config` and update the SELINUX parameter to either `SELINUX=enforcing` or `SELINUX=permissive`:

Example:

```
SELINUX=enforcing
```

2. Update the SELinux running state to either `enforcing` or `permissive`:

If SELinux is currently disabled:

Run the following command to create `/.autorelabel`:

```
# touch /.autorelabel
```

Run the following command to reboot the system:

```
# reboot
```

WARNING: The `autorelabel` may be a long running process and require significant downtime. This should be planned for before executing the reboot!

If SELinux is not currently disabled:

Run one of the following commands to set SELinux's running mode:

- Run the following command to set SELinux running mode to `Enforcing`:

```
# setenforce 1
```

- OR -

- Run the following command to set SELinux running mode to `Permissive`:

```
# setenforce 0
```

References:

1. GRID: NIX-00000096
2. NIST SP 800-53 :: SC-3
3. NIST SP 800-53A :: SC-3.1 (ii)
4. https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/selinux_users_and_administrators_guide/sect-security-enhanced_linux-introduction-selinux_modes

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.	●	●	●

1.3.1.5 Ensure the SELinux mode is enforcing (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

SELinux can run in one of three modes: disabled, permissive, or enforcing:

- `Enforcing` - Is the default, and recommended, mode of operation; in enforcing mode SELinux operates normally, enforcing the loaded security policy on the entire system.
- `Permissive` - The system acts as if SELinux is enforcing the loaded security policy, including labeling objects and emitting access denial entries in the logs, but it does not actually deny any operations. While not recommended for production systems, permissive mode can be helpful for SELinux policy development.
- `Disabled` - Is strongly discouraged; not only does the system avoid enforcing the SELinux policy, it also avoids labeling any persistent objects such as files, making it difficult to enable SELinux in the future

Note: You can set individual domains to permissive mode while the system runs in enforcing mode. For example, to make the httpd_t domain permissive:

```
# semanage permissive -a httpd_t
```

Rationale:

Running SELinux in disabled mode the system not only avoids enforcing the SELinux policy, it also avoids labeling any persistent objects such as files, making it difficult to enable SELinux in the future.

Running SELinux in Permissive mode, though helpful for developing SELinux policy, only logs access denial entries, but does not deny any operations.

Impact:

Running SELinux in Enforcing mode may block intended access to files or processes if the SELinux policy is not correctly configured. If this occurs, review the system logs for details and update labels or policy as appropriate.

Audit:

1. Run the following command to verify SELinux's configured mode:

```
# grep -Pi -- '^\h*SELINUX=enforcing\b' /etc/selinux/config

SELINUX=enforcing
```

2. Run the following command to verify SELinux's current running mode:

```
# getenforce

Enforcing
```

Remediation:

1. Edit `/etc/selinux/config` and update the SELINUX parameter to `SELINUX=enforcing`

Example:

```
SELINUX=enforcing
```

2. Update the SELinux running state to `enforcing`:

- **If SELinux is currently disabled:**

Run the following command to create `/.autorelabel`:

```
# touch /.autorelabel
```

Run the following command to reboot the system:

```
# reboot
```

WARNING: The `autorelabel` may be a long running process and require significant downtime. This should be planned for before executing the reboot!

- **If SELinux is not currently disabled:**

Run the following commands to set SELinux's running mode to `Enforcing`:

```
# setenforce 1
```

References:

1. GRID: NIX-00000095
2. CCI-001084 | CCI-002696
3. NIST SP 800-53 :: SC-3 | NIST SP 800-53 :: SI-6 a
4. NIST SP 800-53A :: SC-3.1 (ii)
5. STIG ID: RHEL-08-010170 | RULE ID: SV-230240r1017059 | CAT II
6. STIG ID: OL08-00-010170 | RULE ID: SV-248548r958518 | CAT II
7. STIG ID: RHEL-09-431010 | RULE ID: SV-258078r958944 | CAT I
8. STIG ID: ALMA-09-041930 | RULE ID: SV-269430r1050313 | CAT II
9. STIG ID: OL09-00-000060 | RULE ID: SV-271452r1091068 | CAT I
10. https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/selinux_users_and_administrators_guide/sect-security-enhanced_linux-introduction-selinux_modes

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.	●	●	●

1.3.1.6 Ensure no unconfined services exist (Manual)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Unconfined processes run in unconfined domains

Rationale:

For unconfined processes, SELinux policy rules are applied, but policy rules exist that allow processes running in unconfined domains almost all access. Processes running in unconfined domains fall back to using DAC rules exclusively. If an unconfined process is compromised, SELinux does not prevent an attacker from gaining access to system resources and data, but of course, DAC rules are still used. SELinux is a security enhancement on top of DAC rules – it does not replace them

Impact:

Confining a service that inherently requires unconfined access to function may disrupt its intended operations. This restriction may lead to downtime, degraded performance, or loss in functionality. It is crucial to analyze and adjust SELinux policies in accordance with site security policies and operational requirements.

Audit:

Run the following command to list any processes running in an unconfined service domain:

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

Ideally, nothing is returned. Any process listed is acceptable **only** if the service genuinely requires unconfined operation and the exception has been analyzed and documented per site policy (for example, some backup or centralized-management daemons); otherwise it should be confined with an appropriate SELinux policy (see the remediation).

Remediation:

Investigate any unconfined processes found during the audit action. If necessary create a customize SELinux policy to allow necessary actions for the service.

Warning: Knowledge about creating and configuring SELinux policies is needed. A Basic example on how to create a policy is included below.

1. Identify the unconfined service: determine the name and process of the service
2. Identify the functionality: determine if the functionality is required for operations
3. Create or add to the custom allow list in the SELinux policy configuration

Example SELinux policy configuration: service_allowlist_policy.te

```
# Example SELinux policy configuration for allowing access to specific
actions and resources for a service

module my_service 1.0;

require {
    type my_service_t;
    type system_resource_t;
    class file { read write execute };
    class dir { read write add_name };
    class tcp_socket name_connect;
}

allow my_service_t system_resource_t:file { read write execute }; # Allow
my_service_t to read, write, and execute files with the system_resource_t
context

allow my_service_t system_resource_t:dir { read write add_name }; # Allow
my_service_t to read and write to directories with the system_resource_t
context

allow my_service_t system_resource_t:tcp_socket name_connect; # Allow
my_service_t to establish TCP connections
```

4. Compile the policy

```
# checkmodule -M -m -o service_allowlist_policy.mod
service_allowlist_policy.te
```

5. Create the package

```
# semodule_package -o service_allowlist_policy.pp -m
service_allowlist_policy.mod
```

6. Load the policy

```
# semodule -i service_allowlist_policy.pp
```

7. Apply the policy to the service

```
# chcon -t my_service_exec_t /path/to/service_binary
```

References:

1. GRID: NIX-00000089
2. https://nb.fedorapeople.org/cvsfedora/web/html/docs/selinux-guide/f10/en-US/sect-Security-Enhanced_Linux-Targeted_Policy-Unconfined_Processes.html
3. https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html/using_selinux/managing-confined-and-unconfined-users_using-selinux
4. https://access.redhat.com/documentation/id-id/red_hat_enterprise_linux/9/pdf/using_selinux/red_hat_enterprise_linux-9-using_selinux-en-us.pdf
5. NIST SP 800-53 Rev. 5: AC-3, MP-2

Additional Information:

Occasionally certain daemons such as backup or centralized management software may require running unconfined. Any such software should be carefully analyzed and documented before such an exception is made.

CIS Controls:

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

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

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `mcstransd` daemon provides category label information to client processes requesting information. The label translations are defined in `/etc/selinux/targeted/setrans.conf`

Rationale:

Since this service is not used very often, remove it to reduce the amount of potentially vulnerable code running on the system.

Audit:

Run the following command and verify `mcstrans` is not installed.

```
# rpm -q mcstrans

package mcstrans is not installed
```

Remediation:

Run the following command to uninstall `mcstrans`:

```
# dnf remove mcstrans
```

References:

1. GRID: NIX-00000094
2. NIST SP 800-53 Rev. 5: AC-3, MP-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.		●	●

1.3.1.8 Ensure SETroubleshoot is not installed (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The SETroubleshoot service notifies desktop users of SELinux denials through a user-friendly interface. The service provides important information around configuration errors, unauthorized intrusions, and other potential errors.

Rationale:

The SETroubleshoot service is an unnecessary daemon to have running on a server, especially if X Windows is disabled.

Audit:

Verify `setroubleshoot` is not installed.

Run the following command:

```
# rpm -q setroubleshoot

package setroubleshoot is not installed
```

Remediation:

Run the following command to uninstall `setroubleshoot`:

```
# dnf remove setroubleshoot
```

References:

1. GRID: NIX-00000093
2. NIST SP 800-53 Rev. 5: AC-3, MP-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	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.	●	●	●

1.4 Configure Bootloader

This section covers securing the bootloader - setting a bootloader password, restricting access to its configuration, and requiring authentication for recovery and emergency boot paths.

The bootloader (GRUB2 on this platform) is the first software the system runs after firmware hand-off. It selects the kernel to boot and the parameters passed to it, so anyone who can interact with it can change how the operating system starts — for example booting to a single-user/recovery shell, disabling Mandatory Access Control (`selinux=0`), or appending kernel arguments — **before any in-OS access control takes effect**. A user with console or physical access is therefore the primary threat this section addresses.

The recommendations reduce that exposure in two ways:

- **Authentication** — require a bootloader superuser password so boot entries cannot be edited (`e`) and the GRUB command line (`c`) cannot be reached without credentials.
- **File protection** — restrict ownership and permissions on the GRUB configuration files so non-root users cannot read the boot configuration (including any stored password hash) or alter boot behavior.

On this platform the GRUB configuration lives under `/boot/grub2/`; on UEFI systems some files also reside under `/boot/efi/EFI/<distro>/`, a vFAT filesystem whose permissions are governed by mount options (`fmask`/`uid`/`gid` in `/etc/fstab`) rather than `chmod`.

Cloud and virtualized environments: The threat this section addresses is interactive access to the boot process. On a Cloud Service Provider (AWS, Azure, GCP, etc.) the tenant has no physical access, provider console access is mediated by the cloud control plane (IAM), and an actor holding control-plane access can reach the guest disk directly (snapshot it, or attach the volume to another instance) — bypassing the bootloader altogether. In such environments a **bootloader password** adds little to the realistic threat model and can complicate provider rescue/recovery workflows, so assess its applicability and record a documented exception per site policy if it is not enforced. **Restricting permissions on the bootloader configuration files** still applies everywhere, as it is local file-permission hardening with no operational downside.

Early boot (initramfs): On a failed boot — or when `rd.break`/`rd.shell` is passed on the kernel command line — the dracut initramfs can drop to a root maintenance shell before the root filesystem is switched. The bootloader password prevents an attacker from adding those arguments interactively, and modern `systemd` refuses the emergency shell when the root account is locked or has a password set. For higher-assurance environments that do not need an interactive early-boot rescue, the dracut shell can additionally be disabled by adding `rd.shell=0 rd.emergency=halt` to the kernel arguments — note this halts on an early-boot failure instead of offering a rescue shell, so recovery then requires boot media.

These recommendations are written for the GRUB2 bootloader. If LILO or another bootloader is in use, enact the equivalent settings for that bootloader.

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 SELinux at boot time).

Impact:

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

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

Audit:

Verify a GRUB bootloader password is set.

Run the following to check the GRUB `user.cfg` file(s) for a password hash:

```
# find /boot -type f -name 'user.cfg' ! -empty -exec grep -HE
'^[[:space:]]*GRUB2?_PASSWORD=' {} +
```

A line similar to the following should be returned:

```
/boot/grub2/user.cfg:GRUB2_PASSWORD=grub.pbkdf2.sha512.10000.<hash>
```

If nothing is returned, a bootloader password is not set.

Remediation:

Create an encrypted password with `grub2-setpassword`:

```
# grub2-setpassword

Enter password: <password>
Confirm password: <password>
```

References:

1. GRID: NIX-00000098
2. NIST SP 800-53 Rev. 5: AC-3
3. NIST SP 800-53A :: AC-3.1
4. STIG ID: RHEL-08-010140 | RULE ID: SV-230234r1017053 | CAT I
5. STIG ID: RHEL-08-010150 | RULE ID: SV-230235r1017054 | CAT I
6. STIG ID: RHEL-09-212010 | RULE ID: SV-257787r1044836 | CAT II

Additional Information:

`grub2-setpassword` outputs the `user.cfg` file which contains the hashed GRUB bootloader password. This utility only supports configurations where there is a single root user.

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.	●	●	●

1.4.2 Ensure access to bootloader config is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

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

Rationale:

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

Impact:

Restricting access to the bootloader configuration files has no impact on normal system operation. Note that permissions on the `/boot/efi` vFat filesystem are governed by filesystem mount options (`fmask`, `uid`, `gid` in `/etc/fstab`) rather than `chmod`, so applying them there requires an `/etc/fstab` change and may require a reboot to take effect.

Audit:

Verify the grub configuration files are appropriately secured.

Run the following to list the mode, owner, and group of the grub configuration files:

```
# find /boot -type f \( -name 'grub*' -o -name 'user.cfg' \) -exec stat -Lc
'%n %a %U %G' {} +
```

Example output:

```
/boot/grub2/grub.cfg 600 root root
/boot/grub2/grubenv 600 root root
/boot/efi/EFI/fedora/grubx64.efi 700 root root
```

For each file returned, verify:

- Files under `/boot/grub2/` are mode `0600` or more restrictive
- Files under `/boot/efi/EFI/` (a vFat filesystem) are mode `0700` or more restrictive
- The owner is user `root`
- The group owner is group `root`

Remediation:

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

- IF - Files are located in `/boot/efi/EFI/*` **Note:** This is a vFat filesystem

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

Example:

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

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

- OR -

- IF - Files are located in `/boot/grub2/*`:

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

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

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

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

References:

1. GRID: NIX-00000097
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. STIG ID: RHEL-09-212025 | RULE ID: SV-257790r991589 | CAT II
5. STIG ID: RHEL-09-212030 | RULE ID: SV-257791r991589 | 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.	●	●	●
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.	●	●	●

1.4.3 Ensure authentication is required when booting into rescue mode (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Rescue mode (single-user mode) boots the system into a minimal environment for administrative maintenance and provides a root shell. When `rescue.service` is configured to run `sulogin` (via `systemd-sulogin-shell`), the root password must be entered before the rescue shell is granted.

Rationale:

Rescue mode grants full root access to the system. If authentication is not required to enter it, anyone able to boot the host into rescue mode is granted root access to all files on the system without credentials. Requiring the root password (via `sulogin`) before the maintenance shell is presented ensures only authorized administrators can use rescue mode.

Impact:

An administrator entering rescue mode from the console will be prompted for the root password. This control requires the root account to have a valid password and is therefore **incompatible with locking the root account** — on a system with a locked or empty root password, `sulogin` will refuse rescue-mode access and recovery will require boot media (or provider/console recovery). Sites that lock the root account should document a deviation or rely on boot-media recovery.

Audit:

Run the following command to verify that rescue mode requires authentication:

```
# systemctl show -p ExecStart --value rescue.service
```

The returned `ExecStart` should invoke `sulogin` (for example via `systemd-sulogin-shell`). *Example output:*

```
{ path=/usr/lib/systemd/systemd-sulogin-shell ;
argv[]=/usr/lib/systemd/systemd-sulogin-shell rescue ; ignore_errors=yes ;
... }
```

If the returned `ExecStart` does not invoke `sulogin`, authentication is not required for rescue mode.

Remediation:

Ensure the root account has a valid password set, then require authentication for rescue mode by creating a systemd drop-in override.

Create or edit `/etc/systemd/system/rescue.service.d/00-require-auth.conf` with the following content:

```
[Service]
ExecStart=-/usr/lib/systemd/systemd-sulogin-shell rescue
```

Reload the systemd manager configuration:

```
# systemctl daemon-reload
```

Default Value:

By default, `rescue.service` invokes `/usr/lib/systemd/systemd-sulogin-shell`, which requires the root password.

Additional Information:

This recommendation is written for the `systemd` init system. Preventing edits to the boot command line (see the bootloader password recommendation) blocks the related vectors of adding `rd.break`, `init=/bin/bash`, or `systemd.unit=emergency.target`; this recommendation ensures the rescue maintenance shell itself requires authentication.

1.4.4 Ensure authentication is required when booting into emergency mode (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

Emergency mode boots the system into the most minimal environment — before most filesystems are mounted — and provides a root shell. When `emergency.service` is configured to run `sulogin` (via `systemd-sulogin-shell`), the root password must be entered before the emergency shell is granted.

Rationale:

Emergency mode grants full root access to the system. If authentication is not required to enter it, anyone able to boot the host into emergency mode is granted root access to all files on the system without credentials. Requiring the root password (via `sulogin`) before the maintenance shell is presented ensures only authorized administrators can use emergency mode.

Impact:

An administrator entering emergency mode from the console will be prompted for the root password. This control requires the root account to have a valid password and is therefore **incompatible with locking the root account** — on a system with a locked or empty root password, `sulogin` will refuse emergency-mode access and recovery will require boot media (or provider/console recovery). Sites that lock the root account should document a deviation or rely on boot-media recovery.

Audit:

Run the following command to verify that emergency mode requires authentication:

```
# systemctl show -p ExecStart --value emergency.service
```

The returned `ExecStart` should invoke `sulogin` (for example via `systemd-sulogin-shell`). *Example output:*

```
{ path=/usr/lib/systemd/systemd-sulogin-shell ;
argv[]=/usr/lib/systemd/systemd-sulogin-shell emergency ; ignore_errors=yes ;
... }
```

If the returned `ExecStart` does not invoke `sulogin`, authentication is not required for emergency mode.

Remediation:

Ensure the root account has a valid password set, then require authentication for emergency mode by creating a systemd drop-in override.

Create or edit `/etc/systemd/system/emergency.service.d/00-require-auth.conf` with the following content:

```
[Service]
ExecStart=-/usr/lib/systemd/systemd-sulogin-shell emergency
```

Reload the systemd manager configuration:

```
# systemctl daemon-reload
```

Default Value:

By default, `emergency.service` invokes `/usr/lib/systemd/systemd-sulogin-shell`, which requires the root password.

Additional Information:

This recommendation is written for the `systemd` init system. Preventing edits to the boot command line (see the bootloader password recommendation) blocks the related vectors of adding `rd.break`, `init=/bin/bash`, or `systemd.unit=emergency.target`; this recommendation ensures the emergency maintenance shell itself requires authentication.

1.4.5 Ensure systemd ability to spawn an interactive boot process is disabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `systemd.confirm_spawn` kernel argument causes `systemd` to prompt interactively at the console before spawning each unit during boot. Removing this argument prevents an interactive boot process.

Rationale:

An interactive boot process lets a person with console access influence which units start during boot, providing an avenue to interfere with the boot process. Ensuring `systemd.confirm_spawn` is not set removes that avenue.

Impact:

None. `systemd.confirm_spawn` is a debugging aid and is not required for normal system operation.

Audit:

Run the following command to verify that `systemd.confirm_spawn` is not set on any boot entry:

```
# grubby --info=ALL | grep -E '^\s*args=' | grep 'systemd.confirm_spawn'
```

Nothing should be returned.

Remediation:

Run the following command to remove the `systemd.confirm_spawn` argument from all boot entries:

```
# grubby --update-kernel ALL --remove-args=systemd.confirm_spawn
```

Then edit `/etc/default/grub` and remove `systemd.confirm_spawn` from the `GRUB_CMDLINE_LINUX` (and `GRUB_CMDLINE_LINUX_DEFAULT`) line so it is not reintroduced on the next kernel or bootloader update. If the system uses `/etc/kernel/cmdline`, remove it there as well.

Default Value:

`systemd.confirm_spawn` is not set by default.

Additional Information:

This recommendation is written for the `systemd` init system.

1.5 Configure Additional Process Hardening

This section covers additional process-hardening Recommendations across two themes:

- **Kernel parameters (`sysctl`)** — filesystem link protections, restriction of kernel information exposure, and address-space layout randomization.
- **Core dump restriction** — limiting or disabling core dumps (via `fs.suid_dumpable`, per-user resource limits in `/etc/security/limits.*`, and `systemd-coredump`), which can otherwise capture sensitive in-memory data such as credentials or keys.

Kernel parameter (`sysctl`) configuration — applies to the `sysctl` Recommendations in this section:

- `sysctl` settings are defined in `.conf` files under `/etc/`, `/run/`, `/usr/local/lib/`, `/usr/lib/`, and `/lib/` — typically in a `sysctl.d` subdirectory. Files must use the `.conf` extension. Vendor defaults usually live under `/usr/lib/` or `/usr/local/lib/`.
- `/usr/lib/systemd/systemd-sysctl --cat-config` lists the effective kernel parameters and the file each is configured in. Within that output, a value listed later overrides the same setting listed earlier, so the **last** occurrence of a parameter is the effective (running) value.
- **`--system` precedence.** When run with `--system`, `sysctl` reads files from the directories below in order (top to bottom); once a filename is loaded, a same-named file in a later directory 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. A file may be fully replaced (a same-named file in a higher-priority directory) or partially replaced (a file ordered later).

- **Setting and applying a value.** Place site settings in a `.conf` file under `/etc/sysctl.d/` (the highest-precedence directory, so it reliably overrides vendor defaults), then apply them to the running system with `sysctl --system`. If more than one file sets the same parameter, the effective value is the one that wins by the precedence above. If that effective value is correct but is set in a file **outside** `/etc/sysctl.d/`, add a matching entry under `/etc/sysctl.d/` so a package update cannot silently change it, and comment out or correct any conflicting values in other `/etc/sysctl.d/` files.

Core dump configuration — applies to the core-dump Recommendations:

- `systemd-coredump` is configured in `/etc/systemd/coredump.conf` and drop-ins under `/etc/systemd/coredump.conf.d/*.conf`. View the merged result with `systemd-analyze cat-config systemd/coredump.conf`; where a setting appears more than once, the **last** value shown is the effective one.
- Process core-dump behavior is also governed by `fs.suid_dumpable` (a `sysctl` parameter, above) and by per-user resource limits in `/etc/security/limits.conf` and files under `/etc/security/limits.d/`.

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 hard links to files they do not own or do not have read/write access to.

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`:

```
#!/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)"
   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`.

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`.

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)"
   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
```

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)
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			

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. When enabled, symbolic links are followed only when outside a sticky world-writable directory, when the user identifier (UID) of the symlink and follower match, or when the directory owner matches the symlink's owner.

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:

Applications that legitimately rely on following symbolic links across differing UIDs within world-writable directories (for example, some services or build tooling that use `/tmp`) may stop working. Validate the change in a test environment before enforcing it in production.

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`:

```
#!/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)"
   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"
```

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:

```
#!/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)"
   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
```

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
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			

1.5.3 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:

Processes that have changed privilege levels (such as SUID programs) and execute-only binaries will not produce a core dump. This is the secure default; the only operational effect is that debugging of such programs via core files is limited.

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:

```
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)"
   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"
```

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)"
   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*'"$l_grep"'\s*=\s*([1-9][0-9]*)/s/^/# /' "$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:

```
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
```

3. Run the following command to load all sysctl configuration files:

```
# sysctl --system
```

Default Value:

fs.suid_dumpable = 0

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			

1.5.4 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:

Setting `kernel.dmesg_restrict` to `1` restricts access to the kernel message buffer (`dmesg`) to users holding the `CAP_SYSLOG` capability; unprivileged users can no longer read it.

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`:

```
#!/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)"
   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 = 1" is set in:
    "/etc/sysctl.d/60-kernel_sysctl.conf"
```

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`.

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)"
   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
```

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
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.	●	●	●

1.5.5 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.

Impact:

With `kernel.kptr_restrict` set to `1` or `2`, kernel pointer addresses exposed through `/proc` (for example `/proc/kallsyms`) and similar interfaces are hidden or zeroed for non-privileged users. Some debugging or performance-profiling tools that read kernel symbol addresses may be affected.

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:

```
#!/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)"
   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 = 1" is set in:
    "/etc/sysctl.d/60-kernel_sysctl.conf"
```

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)"
   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 < <("$l_systemdsysctl" --cat-config | tac | \
   grep -Pio '^\h*#\h*\/[^#\n\r\h]+\.conf\b')
   for l_file in "${a_files[@]}"; do
      grep -Poi '\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:

```
# sysctl --system
```

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
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.		●	●

1.5.6 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.

Impact:

Enabling full address space layout randomization (`kernel.randomize_va_space = 2`) has minimal operational impact. In rare cases, legacy applications or debuggers that assume fixed memory addresses may be affected.

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`:

```
#!/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)"
   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 = 2" is set in:
    "/etc/sysctl.d/60-kernel_sysctl.conf"
```

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`.

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)"
   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
```

3. Run the following command to load all system configuration files:

```
# sysctl --system
```

Default Value:

kernel.randomize_va_space = 2

References:

1. GRID: NIX-00000106
2. NIST SP 800-53 :: SI-16
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.		●	●

1.5.7 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.

More information about the kernel parameter configuration files, their location, and load preference is available in the section overview.

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.

An example is CVE-2026-46333, a logic flaw in the kernel's `__ptrace_may_access()` function that allows an unprivileged local user to copy open file descriptors out of exiting privileged (SUID) processes — such as `ssh-keysign` or `chage` — to disclose sensitive files or escalate to root. Setting `kernel.yama.ptrace_scope` to `2` or `3` requires `CAP_SYS_PTRACE` for `PTRACE_ATTACH`, which an unprivileged attacker does not have, and blocks the public exploits for this issue on systems that are not yet patched.

Impact:

Based on value chosen.

- `1` 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` Only processes with CAP_SYS_PTRACE may use ptrace with PTRACE_ATTACH, or through children calling PTRACE_TRACEME.
- `3` No processes may use ptrace with PTRACE_ATTACH nor via PTRACE_TRACEME. Once set, this sysctl value cannot be changed.

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 = 2
```

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="${l_parameter_name//./(\\.|\\/)}" a_output=() a_files=()
   l_systemdsysctl="$(readlink -e /lib/systemd/systemd-sysctl \
   || readlink -e /usr/lib/systemd/systemd-sysctl)"
   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:

- `kernel.yama.ptrace_scope = 1`
- `kernel.yama.ptrace_scope = 2`
- `kernel.yama.ptrace_scope = 3`

Example output:

```
- "kernel.yama.ptrace_scope = 2" is set in:
    "/etc/sysctl.d/60-kernel_sysctl.conf"
```

Remediation:

1. Review all files being used by systemd sysctl and comment out or remove all `kernel.yama.ptrace_scope` lines that are not `kernel.yama.ptrace_scope=1`, `kernel.yama.ptrace_scope=2`, or `kernel.yama.ptrace_scope=3`.

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)"
   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*'"$l_grep"'\s*=\s*(0|4-9|1[0-9]+)/s/^/# /' "$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.yama.ptrace_scope = {N}
```

Example:

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "kernel.yama.ptrace_scope = 2" >> /etc/sysctl.d/60-
kernel_sysctl.conf
```

Note: The examples use the recommended value `2`; a value of `1`, `2`, or `3` is compliant.

3. Run the following command to load all system configuration files:

```
# sysctl --system
```

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: OL08-00-040282 | RULE ID: SV-248890r991589 | CAT II
6. STIG ID: RHEL-09-213080 | RULE ID: SV-257811r1044872 | CAT II
7. STIG ID: ALMA-09-040940 | RULE ID: SV-269424r1050307 | CAT II
8. STIG ID: OL09-00-002410 | RULE ID: SV-271749r1091959 | CAT II
9. CCI-000366

Additional Information:

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

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.8 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.		●	●

1.5.8 Ensure core file size is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

This recommendation enforces a hard limit of `0` on the `core` file size for all users through `/etc/security/limits.conf` (or a file in `/etc/security/limits.d/`), preventing users from raising the limit and generating core dumps.

- `core` - limits the size of core files that may be created.
- `hard` - a hard resource limit; set by the superuser and enforced by the kernel. A user cannot raise their resource limit above a hard value.

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.

Impact:

Users and processes will be unable to generate core dumps (a hard `core` limit of `0`), which may hinder debugging of applications that crash.

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. 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.

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			

1.5.9 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:

Verify `systemd-coredump ProcessSizeMax` is set to `0`.

Run the following command to list each `systemd-coredump` configuration file and its `ProcessSizeMax` setting:

```
# systemd-analyze cat-config systemd/coredump.conf | grep -P
'^(#\h*/|\h*ProcessSizeMax\h*=)'
```

Each configuration file is listed by its `# <path>` header, followed by any `ProcessSizeMax` setting in that file. Because later files override earlier ones, the **last** `ProcessSizeMax` value shown is the effective setting, and it should be `0`.

Example output:

```
# /etc/systemd/coredump.conf
ProcessSizeMax=2G
# /etc/systemd/coredump.conf.d/60-coredump.conf
ProcessSizeMax=0
```

If no `ProcessSizeMax` line is returned, the compiled-in default of `2G` is in effect and `0` is not configured. Make note of the file in which the effective value is set, for use in the Remediation procedure.

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

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			

1.5.10 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:

Verify `systemd-coredump Storage` is set to `none`.

Run the following command to list each `systemd-coredump` configuration file and its `Storage` setting:

```
# systemd-analyze cat-config systemd/coredump.conf | grep -P
'^(#\h*/|\h*Storage\h*=)'
```

Each configuration file is listed by its `# <path>` header, followed by any `Storage` setting in that file. Because later files override earlier ones, the **last** `Storage` value shown is the effective setting, and it should be `none`.

Example output:

```
# /etc/systemd/coredump.conf
Storage=external
# /etc/systemd/coredump.conf.d/60-coredump.conf
Storage=none
```

If no `Storage` line is returned, the compiled-in default of `external` is in effect and `none` is not configured. Make note of the file in which the effective value is set, for use in the Remediation procedure.

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

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			

1.6 Configure system wide crypto policy

This section covers configuring the system-wide cryptographic policy to disable weak algorithms and protocols and enable strong ones.

The system-wide cryptographic policy configures the algorithms and protocols that the core cryptographic back-ends (for example OpenSSL, GnuTLS, NSS, and OpenSSH) will accept, and is managed with the `update-crypto-policies` command. The following files and directories are involved:

- `/etc/crypto-policies/back-ends` - The individual cryptographic back-end configuration files. Usually linked to the configuration shipped in the crypto-policies package unless a configuration from `local.d` is added.
- `/etc/crypto-policies/config` - A file containing the name of the active crypto-policy set on the system.
- `/etc/crypto-policies/local.d` - Additional configuration shipped by other packages or created by the system administrator. Files placed here (named `<back-end>.config`) are appended to the back-end configuration shipped in the crypto-policies package.
- `/usr/share/crypto-policies/policies` - System policy definition files.
- `/usr/share/crypto-policies/policies/modules` - System subpolicy definition files.
- `/etc/crypto-policies/policies` - Custom policy definition files as configured by the system administrator.
- `/etc/crypto-policies/policies/modules` - Custom subpolicy definition files as configured by the system administrator.
- `/usr/share/crypto-policies/<POLICYNAME>` - Pre-generated back-end configurations for policy POLICYNAME.

The crypto policy definition files have a simple syntax following an INI file key = value syntax

Full policy definition files have suffix `.pol`, subpolicy files have suffix `.pmod`. Subpolicies do not have to have values set for all the keys.

The effective configuration of a policy with subpolicies applied is the same as a configuration from a single policy obtained by concatenating the policy and the subpolicies in question.

The policy files shipped in packages are placed in `/usr/share/crypto-policies/policies` and the subpolicies in `/usr/share/crypto-policies/policies/modules`.

Locally configured policy files should be placed in `/etc/crypto-policies/policies` and subpolicies in `/etc/crypto-policies/policies/modules`.

The policy and subpolicy files must have names in upper-case except for the `.pol` and `.pmod` suffix as the update-crypto-policies command always converts the policy name to upper-case before searching for the policy on the filesystem.

The following predefined policies are included:

- `DEFAULT` - The default system-wide cryptographic policy level offers secure settings for current threat models. It allows the TLS 1.2 and 1.3 protocols, as well as the IKEv2 and SSH2 protocols. The RSA keys and Diffie-Hellman parameters are accepted if they are at least 2048 bits long.
- `LEGACY` - This policy ensures maximum compatibility with Red Hat Enterprise Linux 5 and earlier; it is less secure due to an increased attack surface. In addition to the DEFAULT level algorithms and protocols, it includes support for the TLS 1.0 and 1.1 protocols. The algorithms DSA, 3DES, and RC4 are allowed, while RSA keys and Diffie-Hellman parameters are accepted if they are at least 1023 bits long.
- `FUTURE` - A stricter forward-looking security level intended for testing a possible future policy. This policy does not allow the use of SHA-1 in signature algorithms. It allows the TLS 1.2 and 1.3 protocols, as well as the IKEv2 and SSH2 protocols. The RSA keys and Diffie-Hellman parameters are accepted if they are at least 3072 bits long. If your system communicates on the public internet, you might face interoperability problems.
- `FIPS` - A policy level that conforms with the FIPS 140 requirements. The fips-mode-setup tool, which switches the system into FIPS mode, uses this policy internally. Switching to the FIPS policy does not guarantee compliance with the FIPS 140 standard. You also must re-generate all cryptographic keys after you set the system to FIPS mode. This is not possible in many scenarios.

Working with subpolicies

Subpolicies modify a base policy and are combined with it as a colon-separated list, for example `update-crypto-policies --set DEFAULT:NO-SHA1:NO-SSHCBC`. When working with subpolicies:

- The examples in the Recommendations that follow use the included `DEFAULT` base policy; if another base policy is in use per local site policy, replace `DEFAULT` with the name of the active policy.
- A subpolicy that is not included in the `update-crypto-policies --set` command is not applied.
- A subpolicy must exist before it can be applied.
- Custom subpolicy (`.pmod`) filenames must be upper-case (for example `NO-SHA1.pmod`). The `update-crypto-policies` command converts the name to upper-case before searching the filesystem, so a lower-case file will not be found.
- After changing the active policy, reboot the system (or restart the affected services) so that already-running services and applications use the new settings.

- A directive can be applied globally (for example `cipher = ...` or `mac = ...`) or limited to a scope such as SSH using an `@SSH` suffix (for example `cipher@SSH = ...`). The SSH-related Recommendations in this section use the `@SSH` scope so the change affects only SSH; a global directive would apply it to all back-ends. The `etm` directive is implemented only for SSH and should always be used with the `@SSH` scope.

In addition to the predefined policy levels above, the following subpolicy is referenced in this section:

- `PQ` - Available in newer `crypto-policies` releases. Enables post-quantum and hybrid algorithms - hybrid Module-Lattice-Based Key-Encapsulation Mechanism (ML-KEM) key exchange and pure Module-Lattice-Based Digital Signature Algorithm (ML-DSA) signatures - on top of a base policy, for example `DEFAULT:PQ`. Availability depends on the `crypto-policies` version; older releases do not ship the `PQ` module.

Note: Additional information is available in the `crypto-policies`, `update-crypto-policies`, and `fips-mode-setup` man pages.

1.6.1 Ensure system wide crypto policy is not legacy (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

When a system-wide policy is set up, the default behavior of applications will be to follow the policy. Applications will be unable to use algorithms and protocols that do not meet the policy, unless you explicitly request the application to do so.

The system-wide crypto-policies followed by the crypto core components allow consistently deprecating and disabling algorithms system-wide.

The `LEGACY` policy ensures maximum compatibility with version 5 of the operating system and earlier; it is less secure due to an increased attack surface. In addition to the `DEFAULT` level algorithms and protocols, it includes support for the `TLS 1.0` and `1.1` protocols. The algorithms `DSA`, `3DES`, and `RC4` are allowed, while `RSA keys` and `Diffie-Hellman` parameters are accepted if they are at least 1023 bits long.

Rationale:

If the `LEGACY` system-wide crypto policy is selected, it includes support for TLS 1.0, TLS 1.1, and SSH2 protocols or later. The algorithms DSA, 3DES, and RC4 are allowed, while RSA and Diffie-Hellman parameters are accepted if larger than 1023-bits.

These legacy protocols and algorithms can make the system vulnerable to attacks, including those listed in RFC 7457

Impact:

Environments that require compatibility with older insecure protocols may require the use of the less secure `LEGACY` policy level.

Audit:

Run the following command to verify that the system-wide crypto policy is not `LEGACY`

```
# grep -Pi '^\h*LEGACY\b' /etc/crypto-policies/config
```

Verify that no lines are returned

Remediation:

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

```
# update-crypto-policies --set <CRYPTO POLICY>
```

Example:

```
# update-crypto-policies --set DEFAULT
```

Run the following to make the updated system-wide crypto policy active

```
# update-crypto-policies
```

Default Value:

DEFAULT

References:

1. GRID: NIX-00000374
2. https://docs.redhat.com/en/documentation/red_hat_enterprise_linux/10/html/security_hardening/using-system-wide-cryptographic-policies
3. NIST SP 800-53 Rev. 5: SC-8
4. CRYPTO-POLICIES(7)

Additional Information:

- IF - FIPS is required by local site policy:

Use the built in FIPS policy to aid in conformance to the FIPS 140 requirements. This policy is used automatically by a dracut `initramfs` module and a `systemd` service if the kernel is booted with `fips=1` to switch the system into the FIPS 140 mode. This policy provides at least 112-bit security.

- MACs: SHA-256 or better
- Curves: all prime >= 256 bits
- Signature algorithms: with SHA-256 hash or better (no DSA)
- TLS Ciphers: >= 128-bit key, >= 128-bit block (AES, excluding AES-CBC)
- non-TLS Ciphers: same as TLS Ciphers
- Key exchange: ECDHE, DHE (no DHE-DSS, no RSA)

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.		●	●

1.6.2 Ensure system wide crypto policy disables sha1 hash and signature support (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

SHA-1 (Secure Hash Algorithm 1) is a cryptographic hash function that produces a 160-bit hash value, used both for hashing and as the basis of digital signature algorithms. The system-wide crypto policy sets which hash and signature algorithms the operating system's cryptographic backends will use and accept, and controls SHA-1 in three places: the global `hash` and `sign` algorithm sets, per-backend scope overrides (`hash@<scope>` / `sign@<scope>`), and the `sha1_in_certs` setting that governs SHA-1 signatures in X.509 certificates. The policy also enables SHA-1 by default within a bounded set of legacy-interop scopes, such as package-signature verification, so existing artifacts can still be validated.

Rationale:

The SHA-1 hash function has an inherently weak design, and advancing cryptanalysis has made it vulnerable to attacks. The most significant danger for a hash algorithm is when a "collision" which happens when two different pieces of data produce the same hash value occurs. This hashing algorithm has been considered weak since 2005.

Note: The use of SHA-1 with hash-based message authentication codes (HMAC) do not rely on the collision resistance of the corresponding hash function, and therefore the recent attacks on SHA-1 have a significantly lower impact on the use of SHA-1 for HMAC. Because of this, the recommendation does not disable the hmac-sha1 MAC.

Impact:

Disabling SHA-1 hash and signature support may cause interoperability problems with older peers, legacy certificates, or applications that still rely on SHA-1 signatures. HMAC-SHA1 is intentionally left enabled (see the description) to preserve compatibility where the collision weakness does not apply.

Audit:

All of the following must be true, checked against the resolved policy dump `/etc/crypto-policies/state/CURRENT.pol`:

1. Run the following command to check the global hash and signature algorithm sets for SHA1:

```
# grep -Pi '^\h*(hash|sign)\h*=.*SHA-?1\b' /etc/crypto-
policies/state/CURRENT.pol
```

Nothing should be returned.

2. Run the following command to list any per-backend scope that enables SHA1:

```
# grep -Pi '^\h*(hash|sign)@\S+\h*=.*SHA-?1\b' /etc/crypto-
policies/state/CURRENT.pol
```

Typically nothing is returned.

Note: If a scope is returned, it must be one of the operating system's legacy-interop scopes - `RPM` and, on releases that define them, `DNSSec`, `SMIME`, `PKCS12`, and their import and NSS-internal variants - each of which retains SHA1 to validate existing artifacts. Any other scope returned here does not meet this recommendation.

3. Run the following command to verify SHA1 signatures are disallowed in certificates. The returned value should be `0`; if the directive is absent, the default of `0` (disallowed) applies:

```
# grep -Pi '^\h*sha1_in_certs\h*=' /etc/crypto-policies/state/CURRENT.pol
```

Example output:

```
sha1_in_certs = 0
```

Remediation:

1. Run the following command to set the system-wide cryptographic policy so it includes the operating system's provided NO-SHA1 subpolicy, listed together with every other subpolicy that should remain active (a subpolicy that is not listed is not applied):

```
# update-crypto-policies --set
<CRYPTO_POLICY>:<CRYPTO_SUBPOLICY1>:<CRYPTO_SUBPOLICY2>:<CRYPTO_SUBPOLICY3>
```

Example:

```
# update-crypto-policies --set DEFAULT:NO-SHA1:NO-WEAKMAC:NO-SSHCBC:NO-
SSHCHACHA20
```

2. Run the following command to reboot the system so all services load the updated policy:

```
# reboot
```

Note: NO-SHA1 disables SHA-1 in the global directives only; the operating system's legacy-interop scopes (such as RPM) are left at their defaults. If the audit returned SHA-1 enabled in a specific non-exempt backend scope, create an additional subpolicy module (a .pmod file with an upper-case name) containing hash@<scope> = -SHA1 and include it in the --set list above.

References:

1. GRID: NIX-00000376
2. update-crypto-policies(8)
3. Red Hat Enterprise 8 security hardening
4. https://www.redhat.com/en/blog/how-customize-crypto-policies-rhel-82
5. https://access.redhat.com/articles/3642912
6. NIST SP 800-53 Rev. 5: SC-8
7. crypto-policies(7)

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.		●	●

1.6.3 Ensure system wide crypto policy macs are configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Message Authentication Code (MAC) algorithm is a family of cryptographic functions that is parameterized by a symmetric key. Each of the functions can act on input data (called a “message”) of variable length to produce an output value of a specified length. The output value is called the MAC of the input message.

A MAC algorithm can be used to provide data-origin authentication and data-integrity protection

Rationale:

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 tunnel and capture credentials and information.

A MAC algorithm must be computationally infeasible to determine the MAC of a message without knowledge of the key, even if one has already seen the results of using that key to compute the MAC's of other messages.

Impact:

Restricting the MAC algorithms available to the system-wide policy may cause interoperability problems with older clients or servers that only offer the disabled (weaker, 128-bit truncated) MAC algorithms.

Audit:

Run the following command to verify that no weak (less than 128-bit) MAC is enabled in the system-wide crypto policy. Nothing should be returned:

```
# grep -Pi -- '^\h*mac(@[^=\h]+)?\h*=\h*([^#\n\r]+)?-64\b' /etc/crypto-policies/state/CURRENT.pol
```

Note: In the resolved policy, `UMAC-64` is the weak 64-bit MAC targeted here; `UMAC-128` is a strong 128-bit MAC and is expected to remain enabled.

Remediation:

Create or edit a file ending in `.pmod` in `/etc/crypto-policies/policies/modules/` and add the following line to disable weak (less than 128-bit) MACs:

```
mac = -*-64
```

Example:

```
# printf '%s\n' "# Disable weak (less than 128-bit) MACs" "mac = -*-64" >>
/etc/crypto-policies/policies/modules/NO-WEAKMAC.pmod
```

Apply the subpolicy together with every other subpolicy that should remain active (a subpolicy that is not listed is not applied), then reboot:

```
# update-crypto-policies --set
<CRYPTO_POLICY>:<SUBPOLICY1>:<SUBPOLICY2>:<SUBPOLICY3>
```

Example:

```
# update-crypto-policies --set DEFAULT:NO-SHA1:NO-WEAKMAC
```

Reboot the system for the updated policy to take effect.

References:

1. GRID: NIX-00000377
2. crypto-policies(7)
3. update-crypto-policies(8)
4. Red Hat Enterprise 8 security hardening
5. https://www.redhat.com/en/blog/how-customize-crypto-policies-rhel-82
6. https://csrc.nist.gov/glossary/term/message_authentication_code_algorithm
7. NIST SP 800-53 :: CM-6, SC-8, SC-8(1), SC-13
8. https://nvd.nist.gov/vuln/detail/CVE-2008-5161

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.		●	●

1.6.4 Ensure system wide crypto policy disables cbc for ssh (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Cipher Block Chaining (CBC) is an algorithm that uses a block cipher.

Rationale:

A vulnerability exists in SSH messages that employ CBC mode that may allow an attacker to recover plaintext from a block of ciphertext. If exploited, this attack can potentially allow an attacker to recover up to 32 bits of plaintext from an arbitrary block of ciphertext from a connection secured using the SSH protocol.

Impact:

CBC ciphers might be the only common ciphers when connecting to older SSH clients and servers

Audit:

Run the following command to verify if CBC has been disabled system-wide:

```
# grep -Pi -- '^\h*cipher\h*=\h*([^#\n\r]+)?-CBC\b' /etc/crypto-
policies/state/CURRENT.pol
```

Nothing should be returned.

Run the following command to verify CBC is disabled for SSH:

```
# grep -Pi -- '^\h*cipher@(lib|open)ssh(-server|-client)?\h*=\h*([^#\n\r]+)?-
CBC\b' /etc/crypto-policies/state/CURRENT.pol
```

Nothing should be returned.

Remediation:

Create or edit a file in `/etc/crypto-policies/policies/modules/` ending in `.pmod` and add or modify **one** of the following lines:

```
cipher@SSH = -*-CBC # Disables the CBC cipher for SSH
```

Example:

```
# printf '%s\n' "# This is a subpolicy to disable all CBC mode ciphers" "#
for the SSH protocol (libssh and OpenSSH)" "cipher@SSH = -*-CBC" >>
/etc/crypto-policies/policies/modules/NO-SSHCBC.pmod
```

Run the following command to update the system-wide cryptographic policy

```
# update-crypto-policies --set
<CRYPTO_POLICY>:<CRYPTO_SUBPOLICY1>:<CRYPTO_SUBPOLICY2>:<CRYPTO_SUBPOLICY3>
```

Example:

```
# update-crypto-policies --set DEFAULT:NO-SHA1:NO-WEAKMAC:NO-SSHCBC
```

Run the following command to reboot the system to make your cryptographic settings effective for already running services and applications:

```
# reboot
```

References:

1. GRID: NIX-00000378
2. crypto-policies(7)
3. update-crypto-policies(8)
4. Red Hat Enterprise 8 security hardening
5. https://www.redhat.com/en/blog/how-customize-crypto-policies-rhel-82
6. NIST SP 800-53 Rev. 5: SC-8
7. https://nvd.nist.gov/vuln/detail/CVE-2008-5161

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.		●	●

1.6.5 Ensure system wide crypto policy disables chacha20-poly1305 for ssh (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

ChaCha20-Poly1305 is an authenticated encryption with additional data (AEAD) algorithm, that combines the ChaCha20 stream cipher with the Poly1305 message authentication code. Its usage in IETF protocols is standardized in RFC 8439.

Note: If `CVE-2023-48795` (the Terrapin SSH attack) has been addressed — for example, the system runs OpenSSH 9.6 or later, where strict key exchange is negotiated automatically — this recommendation is Not Applicable.

Rationale:

A vulnerability exists in ChaCha20-Poly1305 as referenced in `CVE-2023-48795`

Impact:

Disabling the `chacha20-poly1305` cipher for SSH may cause interoperability problems with clients or servers that only offer that cipher.

Audit:

- IF - `CVE-2023-48795` (the Terrapin attack) has been addressed — for example, the system runs OpenSSH 9.6 or later, where strict key exchange is negotiated automatically — and this meets local site policy, this recommendation is **Not Applicable**.

Run the following command to display the effective SSH cipher configuration:

```
# grep -Pi -- '^\h*cipher(@(lib|open)ssh(-server|-client)?)?\h*='
/etc/crypto-policies/state/CURRENT.pol
```

If a `cipher@...ssh...` line is present it governs SSH; otherwise the global `cipher` line applies. `chacha20-poly1305` should not appear in the governing line.

Example output (compliant):

```
cipher@libssh = AES-256-GCM AES-256-CTR AES-128-GCM AES-128-CTR
cipher@openssh-client = AES-256-GCM AES-256-CTR AES-128-GCM AES-128-CTR
cipher@openssh-server = AES-256-GCM AES-256-CTR AES-128-GCM AES-128-CTR
```

Remediation:

- IF - CVE-2023-48795 (the Terrapin attack) has been addressed — for example, the system runs OpenSSH 9.6 or later — and this meets local site policy, this recommendation may be skipped.

Create or edit a file in /etc/crypto-policies/policies/modules/ ending in .pmod and add or modify **one** of the following lines:

```
cipher@SSH = -CHACHA20-POLY1305 # Disables the chacha20-poly1305 cipher for
SSH
```

Example:

```
# printf '%s\n' "# This is a subpolicy to disable the chacha20-poly1305
ciphers" "# for the SSH protocol (libssh and OpenSSH)" "cipher@SSH = -
CHACHA20-POLY1305" >> /etc/crypto-policies/policies/modules/NO-
SSHCHACHA20.pmod
```

Run the following command to update the system-wide cryptographic policy

```
# update-crypto-policies --set
<CRYPTO_POLICY>:<CRYPTO_SUBPOLICY1>:<CRYPTO_SUBPOLICY2>:<CRYPTO_SUBPOLICY3>
```

Example:

```
# update-crypto-policies --set DEFAULT:NO-SHA1:NO-WEAKMAC:NO-SSHCBC:NO-
SSHCHACHA20
```

Run the following command to reboot the system to make your cryptographic settings effective for already running services and applications:

```
# reboot
```

References:

1. GRID: NIX-00000379
2. crypto-policies(7)
3. update-crypto-policies(8)
4. Red Hat Enterprise 8 security hardening
5. https://www.redhat.com/en/blog/how-customize-crypto-policies-rhel-82
6. NIST SP 800-53 Rev. 5: SC-8
7. https://nvd.nist.gov/vuln/detail/CVE-2023-48795

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.		●	●

1.6.6 Ensure system wide crypto policy disables EtM for ssh (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Encrypt-then-MAC (EtM) - The ciphertext is generated by encrypting the plaintext and then appending a MAC of the encrypted plaintext

Rationale:

There is an effective attack against SSH's use of Cypher-Block-Chaining (CBC) with Encrypt-then-MAC as referenced in CVE-2023-48795

Impact:

Disabling Encrypt-then-MAC (EtM) for SSH may cause interoperability problems with clients or servers that only offer EtM MAC modes.

Audit:

- IF - CVE-2023-48795 (the Terrapin attack) has been addressed — for example, the system runs OpenSSH 9.6 or later, where strict key exchange is negotiated automatically — or if CBC is disabled for the OpenSSH server, this recommendation is Not Applicable.

Run the following command to verify EtM is disabled for SSH:

```
# grep -Psi -- '^\h*etm\b' /etc/crypto-policies/state/CURRENT.pol
```

Verify output includes either:

```
etm@libssh = DISABLE_ETM
etm@openssh-client = DISABLE_ETM
etm@openssh-server = DISABLE_ETM
```

- OR -

```
etm = DISABLE_ETM
```

Note: The ability to disable EtM through system wide crypto policy was added in version 9.3

Remediation:

- IF - CVE-2023-48795 (the Terrapin attack) has been addressed — for example, the system runs OpenSSH 9.6 or later — and this meets local site policy, this recommendation may be skipped.

Create or edit a file in /etc/crypto-policies/policies/modules/ ending in .pmod and add or modify the following line:

```
etm@SSH = DISABLE_ETM # This disables EtM for openSSH and libssh
```

Example:

```
# printf '%s\n' "# This is a subpolicy to disable Encrypt then MAC" "# for
the SSH protocol (libssh and OpenSSH)" "etm@SSH = DISABLE_ETM" >>
/etc/crypto-policies/policies/modules/NO-SSHETM.pmod
```

Run the following command to update the system-wide cryptographic policy

```
# update-crypto-policies --set
<CRYPTO_POLICY>:<CRYPTO_SUBPOLICY1>:<CRYPTO_SUBPOLICY2>:<CRYPTO_SUBPOLICY3>
```

Example:

```
# update-crypto-policies --set DEFAULT:NO-SHA1:NO-WEAKMAC:NO-SSHCBC:NO-
SSHCHACHA20:NO-SSHETM
```

Run the following command to reboot the system to make your cryptographic settings effective for already running services and applications:

```
# reboot
```

References:

1. GRID: NIX-00000380
2. crypto-policies(7)
3. update-crypto-policies(8)
4. Red Hat Enterprise 8 security hardening
5. https://www.redhat.com/en/blog/how-customize-crypto-policies-rhel-82
6. NIST SP 800-53 Rev. 5: SC-8
7. https://access.redhat.com/errata/RHSA-2024:0499
8. https://nvd.nist.gov/vuln/detail/CVE-2023-48795

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.		●	●

1.6.7 Ensure system wide crypto policy enables post-quantum algorithms (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

A system-wide crypto policy subpolicy named `PQ` enables post-quantum and hybrid cryptographic algorithms - hybrid Module-Lattice-Based Key-Encapsulation Mechanism (ML-KEM) key exchange and pure Module-Lattice-Based Digital Signature Algorithm (ML-DSA) signatures. It is applied on top of a base policy, for example `DEFAULT:PQ`.

The `PQ` subpolicy is available in newer `crypto-policies` releases. Where it is available, applying it enables these algorithms with the highest priority for the back-ends that support them (for example OpenSSL, GnuTLS, and NSS for TLS, and, on newer OpenSSH versions, the SSH key exchange).

Rationale:

Cryptographically relevant quantum computers threaten classical public-key algorithms such as RSA, ECDSA, and classical Diffie-Hellman key exchange. An adversary can capture encrypted traffic now and decrypt it once such a computer becomes available - a "harvest now, decrypt later" attack.

Enabling the hybrid post-quantum key exchange provided by the `PQ` subpolicy protects key exchange against this threat. Because the construction is hybrid - a classical algorithm combined with ML-KEM - the connection is no weaker than the classical algorithm alone, even if one component is later found to be weak.

Impact:

The `PQ` subpolicy is only available on `crypto-policies` releases that ship the `PQ` module. On releases that predate it, this recommendation is Not Applicable and no change is possible.

The post-quantum key exchange is hybrid (a classical algorithm combined with ML-KEM), so peers that do not support the post-quantum algorithms negotiate the classical algorithm and interoperability is preserved. Post-quantum SSH key exchange additionally requires a recent OpenSSH version; on older OpenSSH the `PQ` subpolicy affects only the non-SSH back-ends.

Audit:

- IF - the PQ crypto-policy subpolicy is not available on this release - the file /usr/share/crypto-policies/policies/modules/PQ.pmod does not exist - this recommendation is **Not Applicable**.

Run the following command to verify the active system-wide crypto policy includes the PQ subpolicy:

```
# update-crypto-policies --show
```

The output should include the PQ subpolicy appended to the active base policy, for example:

```
DEFAULT:PQ
```

Remediation:

- IF - the PQ subpolicy is not available on this release (the file /usr/share/crypto-policies/policies/modules/PQ.pmod does not exist), this recommendation is Not Applicable and the following steps do not apply.

Run the following command to apply the PQ subpolicy on top of the active base policy. Replace DEFAULT with the active base policy if it differs, and include any other subpolicies already in use:

```
# update-crypto-policies --set DEFAULT:PQ
```

Run the following command to reboot the system to make your cryptographic settings effective for already running services and applications:

```
# reboot
```

Default Value:

The PQ subpolicy is not enabled by default.

References:

1. GRID: NIX-00000413
2. https://docs.redhat.com/en/documentation/red_hat_enterprise_linux/10/html/security_hardening/using-system-wide-cryptographic-policies
3. NIST SP 800-53 Rev. 5: SC-8
4. CRYPTO-POLICIES(7)
5. UPDATE-CRYPTO-POLICIES(8)

Additional Information:

Post-quantum key exchange for OpenSSH requires a recent OpenSSH release (the mlkem768x25519-sha256 key exchange algorithm). On systems with an older OpenSSH, applying the PQ subpolicy affects only the other supported back-ends (for example TLS via OpenSSL, GnuTLS, and NSS).

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.		●	●

1.7 Configure Command Line Warning Banners

This section covers configuring the command-line warning banners (`/etc/motd`, `/etc/issue`, and `/etc/issue.net`) and restricting access to those files.

Presenting a warning message prior to the normal user login may assist in the prosecution of trespassers on the computer system. Changing some of these login banners also has the side effect of hiding OS version information and other detailed system information from attackers attempting to target specific exploits at a system.

Guidelines published by the US Department of Defense require that warning messages include at least the name of the organization that owns the system, the fact that the system is subject to monitoring and that such monitoring is in compliance with local statutes, and that use of the system implies consent to such monitoring. It is important that the organization's legal counsel review the content of all messages before any system modifications are made, as these warning messages are inherently site-specific. More information (including citations of relevant case law) can be found at http://www.justice.gov/criminal/cybercrime/

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

Note: The text provided in the remediation actions for these items is intended as an example only. Please edit to include the specific text for your organization as approved by your legal department.

1.7.1 Ensure /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:

The `/etc/motd` file is displayed after a user authenticates. Displaying operating system and patch-level information — for example the `\m`, `\r`, `\s`, and `\v` escape sequences, or the operating system name — provides detailed system information to attackers attempting to target specific exploits of a system. Where an organization uses the message of the day to convey acceptable-use or monitoring notices, that content should follow local site policy. Authorized users can obtain system information directly (for example by running `uname -a`) once logged in, so it does not need to be presented in the banner.

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/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 any `/etc/motd` or `/etc/motd.d/*` file to remove the `\m`, `\r`, `\s`, and `\v` escape sequences and any reference to the OS platform, and ensure the remaining contents follow local site policy.

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 -

If a message of the day is not in use, remove the 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			

1.7.2 Ensure /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.

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
```

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
```

Remediation:

Edit the `/etc/issue` file with the appropriate contents according to your site policy, remove any instances of `\m` , `\r` , `\s` , `\v` or references to the `OS platform`

Example:

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

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			

1.7.3 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
```

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:

```
# 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			

1.7.4 Ensure access to /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.

Rationale:

- IF - the `/etc/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 `/etc/motd` exists, `Access` is `644` or more restrictive, `Uid` and `Gid` are both `0/root`:

```
# [ -e /etc/motd ] && stat -Lc 'Access: (%#a/%A)  Uid: ( %u/ %U) Gid: ( %g/
%G)' /etc/motd

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 `/etc/motd`:

```
# 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:

```
# rm /etc/motd
```

Default Value:

Access: (0644/-rw-r--r--) Uid: (0/ root) Gid: (0/ root)

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.	●	●	●

1.7.5 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

Access: (0644/-rw-r--r--)  Uid: ( 0/ root) Gid: ( 0/ root)
```

Remediation:

Run the following commands to set mode, owner, and group on `/etc/issue`:

```
# 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

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.	●	●	●

1.7.6 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

Access: (0644/-rw-r--r--)  Uid: ( 0/ root)   Gid: ( 0/ root)
```

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

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.	●	●	●

1.8 Configure GNOME Display Manager

This section covers hardening the GNOME Display Manager (GDM) - the login banner, user list, screen lock, automatic mounting and autorun, and remote display access.

The GNOME Display Manager (GDM) is a program that manages graphical display servers and handles graphical user logins.

GDM and the GNOME desktop read their settings through `dconf` / `GSettings`. System-wide defaults are set in keyfiles under `/etc/dconf/db/local.d/` and are made mandatory (so users cannot override them) by listing the corresponding keys under `/etc/dconf/db/local.d/locks/`; both are compiled into the binary database with `dconf update`. The system database is bound to users through the dconf profile `/etc/dconf/profile/user`, which references `user-db:user` and `system-db:local`.

Two points apply to every recommendation in this section:

- **Locking matters.** Without a corresponding entry under `locks/`, a system-wide value is only a default and a user's own setting takes precedence. Each recommendation therefore verifies both that the value is set system-wide and that the key is locked.
- **Changes require re-login.** `dconf update` must be run, and users must log out and back in, before system-wide settings take effect.

Note: If GDM is not installed on the system, this section is not applicable and should be skipped.

1.8.1 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.

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 users 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:

```
# 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 users 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:

```
# 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.

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			

1.8.2 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:

```
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:

```
#!/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

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.		●	●

1.8.3 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.

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.

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:

```
# 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:

```
#!/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
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.	●	●	●

1.8.4 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.

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:

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:

```
#!/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
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.	●	●	●

1.8.5 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.

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:

```
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:

```
#!/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.	●	●	●

1.8.6 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 comment out the `Enable=true` line in the `[xdmcp]` block:

Example File:

```
# 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.		●	●
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.		●	●

1.8.7 Ensure Wayland is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The GNOME Display Manager (GDM) uses the Wayland display server by default. The `WaylandEnable` setting in the `[daemon]` block of `/etc/gdm/custom.conf` (or `/etc/gdm3/custom.conf`) controls this behavior; when it is set to `false`, GDM and user login sessions fall back to the legacy Xorg (X11) display server.

Rationale:

Wayland isolates graphical clients from one another. Under the legacy X11 protocol any connected client can read another client's keystrokes, capture the contents of other windows, and inject input events, because X11 provides no isolation between clients sharing a display. Wayland enforces this isolation at the compositor, so a compromised or malicious application cannot observe or manipulate other graphical sessions. Leaving Wayland enabled (the default) rather than falling back to Xorg preserves this isolation and reduces the graphical attack surface.

Impact:

Wayland is the default display server, so leaving it enabled has no impact on a default installation. A small number of legacy applications and tools (for example some screen-capture, remote-desktop, or assistive technologies) require an Xorg session; where those are a documented operational requirement, this recommendation may not be applicable. Legacy X11 applications otherwise continue to run under Wayland through Xwayland.

Audit:

Run the following command to verify that GDM is not configured to disable Wayland:

```
# grep -Psi -- '^\h*WaylandEnable\h*=\h*false\b' /etc/gdm/custom.conf
/etc/gdm3/custom.conf
```

Nothing should be returned.

Note: If GDM is not installed on the system, this recommendation is not applicable.

Remediation:

If Wayland has been disabled, edit the GDM configuration file (`/etc/gdm/custom.conf` or `/etc/gdm3/custom.conf`) and remove or comment out the `WaylandEnable=false` line in the `[daemon]` block so that GDM uses the default Wayland display server. For example:

```
# sed -ri '/^\s*\[daemon\]/,/^\s*\[/ s/^(\s*WaylandEnable\s*=\s*false)/#
\1/I' /etc/gdm/custom.conf
```

Reboot the system, or restart GDM, for the change to take effect.

Default Value:

Wayland is enabled by default; `WaylandEnable` is not set to `false`.

References:

1. GRID: NIX-00000419

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 Services

While applying system updates and patches helps correct known vulnerabilities, one of the best ways to protect the system against as yet unreported vulnerabilities is to disable all services that are not required for normal system operation. This prevents the exploitation of vulnerabilities discovered at a later date. If a service is not enabled, it cannot be exploited. The actions in this section of the document provide guidance on some services which can be safely disabled and under which circumstances, greatly reducing the number of possible threats to the resulting system. Additionally, some services which should remain enabled but with secure configuration are covered as well as insecure service clients.

2.1 Configure Server Services

This section covers server services that may be installed on the system. Many are installed by default or pulled in to satisfy a dependency, and every running service expands the system's attack surface. Any service that is not required for the system's intended role should not be in use.

A service is **not in use** when its package is removed, or - where the package must remain - when its service and socket are both stopped and masked. Removing the package is preferred, as it also removes the associated binaries and configuration.

- 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
```

Masking takes effect even when the corresponding package is not installed, and the service and socket will remain masked; if the package is later installed to meet an approved requirement, its service and/or socket must be unmasked before it can be enabled or started.

Note: This should not be considered a comprehensive list of services not required for normal system operation. You may wish to consider additions to those listed here for your environment.

2.1.1 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 system even if they lacked permissions to mount it themselves. Requiring an administrator to manually mount media can be part of a defense-in-depth strategy that reduces the risk of unapproved software or information being introduced, or proprietary software or information being exfiltrated.

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:

```
# rpm -q autofs

package autofs is not installed
```

- OR -

- IF - the package is required for dependencies:

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

```
# 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 `autofs` package:

```
# systemctl stop autofs.service
# dnf remove 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. NIST SP 800-53 Revision 5 :: IA-3
4. NIST SP 800-53A :: IA-3.1 (ii)
5. RHEL 8 STIG Vul ID: V-230502
6. RHEL 8 STIG Rule ID: SV-230502r627750

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.	●	●	●

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 the avahi package is not installed:

```
# rpm -q avahi

package avahi is not installed
```

- 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:

```
# 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 package:

```
# systemctl stop avahi-daemon.socket avahi-daemon.service
# dnf remove avahi
```

- OR -

- IF - the avahi 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.		●	●
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.		●	●

2.1.3 Ensure cockpit web services are not in use (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

cockpit Web services provide the ability to remotely administer the system via a web browser.

Rationale:

Unless there is a local site approved requirement for remote web based administration on the system, cockpit web server packages should be removed to reduce the potential attack surface.

Impact:

Removal of cockpit web service packages will remove the ability for remote web administration of that system.

- IF - the cockpit web services package is required for a dependency, any related service or socket should be stopped and masked.

Audit:

Run the following command to verify `cockpit-ws` is not installed:

```
# rpm -q cockpit-ws

package cockpit-ws is not installed
```

- OR -

- IF - a package is installed **and** is required for dependencies:

Run the following command to verify `cockpit.socket`, and `cockpit.service` are not enabled:

```
# systemctl is-enabled cockpit.socket cockpit.service 2>/dev/null | grep
'enabled'

Nothing should be returned
```

Run the following command to verify `cockpit.socket`, and `cockpit.service` are not active:

```
# systemctl is-active cockpit.socket cockpit.service 2>/dev/null | grep
'^active'

Nothing should be returned
```

Remediation:

Run the following commands to stop `cockpit.socket`, and `cockpit.service`, and remove the `cockpit-ws` package:

```
# systemctl stop cockpit.socket cockpit.service
# dnf remove cockpit-ws
```

- OR -

- IF - a package is installed **and** is required for dependencies:

Run the following commands to stop and mask `cockpit.socket`, and `cockpit.service`:

```
# systemctl stop cockpit.socket cockpit.service
# systemctl mask cockpit.socket cockpit.service
```

References:

1. GRID: NIX-00000126
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.		●	●

2.1.4 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. There are two versions of the DHCP protocol `DHCPv4` and `DHCPv6`. At startup the server may be started for one or the other via the `-4` or `-6` arguments.

Rationale:

Unless a system is specifically set up to act as a DHCP server, it is recommended that the `dhcp-server` package be removed to reduce the potential attack surface.

Impact:

There may be packages that are dependent on the `dhcp-server` package. If the `dhcp-server` package is removed, these dependent packages will be removed as well. Before removing the `dhcp-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 `dhcpd.service` and `dhcpd6.service` leaving the `dhcp-server` package installed.

Audit:

Run the following command to verify `dhcp-server` is not installed:

```
# rpm -q dhcp-server

package dhcp-server is not installed
```

- OR -

- IF - the package is required for dependencies:

Run the following command to verify `dhcpd.service` and `dhcpd6.service` are not enabled:

```
# systemctl is-enabled dhcpd.service dhcpd6.service 2>/dev/null | grep
'enabled'

Nothing should be returned
```

Run the following command to verify `dhcpd.service` and `dhcpd6.service` are not active:

```
# systemctl is-active dhcpd.service dhcpd6.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 `dhcpd.service` and `dhcpd6.service` and remove the `dhcp-server` package:

```
# systemctl stop dhcpd.service dhcpd6.service
# dnf remove dhcp-server
```

- OR -

- IF - the `dhcp-server` package is required as a dependency:

Run the following commands to stop and mask `dhcpd.service` and `dhcpd6.service`:

```
# systemctl stop dhcpd.service dhcpd6.service
# systemctl mask dhcpd.service dhcpd6.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.		●	●

2.1.5 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.

Rationale:

Unless a system is specifically designated to act as a DNS server, it is recommended that the package be removed to reduce the potential attack surface.

Impact:

There may be packages that are dependent on the `bind` package. If the `bind` package is removed, these dependent packages will be removed as well. Before removing the `bind` package, review any dependent packages to determine if they are required on the system.

- IF - a dependent package is required: stop and mask the `named.service` leaving the `bind` package installed.

Audit:

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

```
# rpm -q bind

package bind is not installed
```

- 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

- 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 `bind` package:

```
# systemctl stop named.service
# dnf remove bind
```

- OR -

- IF - the `bind` package is required as a dependency:

Run the following commands to stop and mask `named.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.		●	●

2.1.6 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 the following command to verify `dnsmasq` is not installed:

```
# rpm -q dnsmasq

package dnsmasq is not installed
```

- 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

- 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
# dnf remove 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.		●	●

2.1.7 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 package can be removed 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 `smb.service` leaving the `samba` package installed.

Audit:

Run the following command to verify `samba` package is not installed:

```
# rpm -q samba

package samba is not installed
```

- OR -

- IF - the package is required for dependencies:

Run the following command to verify `smb.service` is not enabled:

```
# systemctl is-enabled smb.service 2>/dev/null | grep 'enabled'

Nothing should be returned
```

Run the following command to verify the `smb.service` is not active:

```
# systemctl is-active smb.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 `smb.service` and remove `samba` package:

```
# systemctl stop smb.service
# dnf remove samba
```

- OR -

- IF - the `samba` package is required as a dependency:

Run the following commands to stop and mask the `smb.service`:

```
# systemctl stop smb.service
# systemctl mask smb.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.		●	●

2.1.8 Ensure ftp server services are not in use (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

FTP (File Transfer Protocol) is a traditional and widely used standard tool for transferring files between a server and clients over a network, especially where no authentication is necessary (permits anonymous users to connect to a server).

Rationale:

Unless there is a need to run the system as a FTP server, it is recommended that the package be removed to reduce the potential attack surface.

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 `vsftpd` is not installed:

```
# rpm -q vsftpd

package vsftpd is not installed
```

- 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
```

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 `vsftpd` package:

```
# systemctl stop vsftpd.service
# dnf remove 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
3. NIST SP 800-53A :: CM-6.1 (iv)
4. RHEL 8 STIG Vul ID: V-230558
5. RHEL 8 STIG Rule ID: SV-230558r627750

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.		●	●

2.1.9 Ensure message access server services are not in use (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`dovecot` and `cyrus-imapd` are open source IMAP and POP3 server packages 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.

Impact:

There may be packages that are dependent on `dovecot` and `cyrus-imapd` packages. If `dovecot` and `cyrus-imapd` packages are removed, these dependent packages will be removed as well. Before removing `dovecot` and `cyrus-imapd` 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`, `dovecot.service` and `cyrus-imapd.service` leaving `dovecot` and `cyrus-imapd` packages installed.

Audit:

Run the following command to verify `dovecot` and `cyrus-imapd` are not installed:

```
# rpm -q dovecot cyrus-imapd

package dovecot is not installed
package cyrus-imapd is not installed
```

- OR -

- IF - a package is installed **and** is required for dependencies:

Run the following commands to verify `dovecot.socket`, `dovecot.service`, and `cyrus-imapd.service` are not enabled:

```
# systemctl is-enabled dovecot.socket dovecot.service cyrus-imapd.service
2>/dev/null | grep 'enabled'

Nothing should be returned
```

Run the following command to verify `dovecot.socket`, `dovecot.service`, and `cyrus-imapd.service` are not active:

```
# systemctl is-active dovecot.socket dovecot.service cyrus-imapd.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 `dovecot.socket`, `dovecot.service`, and `cyrus-imapd.service`, and remove `dovecot` and `cyrus-imapd` packages:

```
# systemctl stop dovecot.socket dovecot.service cyrus-imapd.service
# dnf remove dovecot cyrus-imapd
```

- OR -

- IF - a package is installed **and** is required for dependencies:

Run the following commands to stop and mask `dovecot.socket`, `dovecot.service`, and `cyrus-imapd.service`:

```
# systemctl stop dovecot.socket dovecot.service cyrus-imapd.service
# systemctl mask dovecot.socket dovecot.service cyrus-imapd.service
```

References:

1. GRID: NIX-00000134
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.		●	●

2.1.10 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 require access to network shares or the ability to provide network file system services for other host's network shares, it is recommended that the `nfs-utils` package be removed to reduce the attack surface of the system.

Impact:

Many of the `libvirt` packages used by Enterprise Linux virtualization are dependent on the `nfs-utils` package. If the `nfs-utils` package is removed, these dependent packages will be removed as well. Before removing the `nfs-utils` 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-utils` package installed.

Audit:

Run the following command to verify `nfs-utils` is not installed:

```
# rpm -q nfs-utils

package nfs-utils is not installed
```

- 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:

```
# 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-utils` package:

```
# systemctl stop nfs-server.service
# dnf remove nfs-utils
```

- OR -

- IF - the `nfs-utils` package is required as a dependency:

Run the following commands to stop and mask the `nfs-server.service`:

```
# systemctl stop nfs-server.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.		●	●

2.1.11 Ensure print server services are not in use (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The Common Unix Print System (CUPS) provides the ability to print to both local and network printers. A system running CUPS can also accept print jobs from remote systems and print them to local printers. It also provides a web based remote administration capability.

Rationale:

If the system does not need to print jobs or accept print jobs from other systems, it is recommended that CUPS be removed to reduce the potential attack surface.

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:

```
# rpm -q cups

package cups is not installed
```

- 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:

```
# 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
# dnf remove 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.		●	●
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.		●	●

2.1.12 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:

```
# rpm -q rpcbind

package rpcbind is not installed
```

- OR -

- IF - the `rpcbind` package is required as a dependency:

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
# dnf remove 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

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.		●	●

2.1.13 Ensure rsync services are not in use (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

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

Rationale:

Unless required, the `rsync-daemon` package should be removed to reduce the potential attack surface.

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

Impact:

There may be packages that are dependent on the `rsync-daemon` package. If the `rsync-daemon` package is removed, these dependent packages will be removed as well. Before removing the `rsync-daemon` package, review any dependent packages to determine if they are required on the system.

- IF - a dependent package is required: stop and mask the `rsyncd.socket` and `rsyncd.service` leaving the `rsync-daemon` package installed.

Audit:

Run the following command to verify the `rsync-daemon` package is not installed:

```
# rpm -q rsync-daemon

package rsync-daemon is not installed
```

- OR -

- IF - the `rsync-daemon` package is required as a dependency:

Run the following command to verify `rsyncd.socket` and `rsyncd.service` are not enabled:

```
# systemctl is-enabled rsyncd.socket rsyncd.service 2>/dev/null | grep
'enabled'

Nothing should be returned
```

Run the following command to verify `rsyncd.socket` and `rsyncd.service` are not active:

```
# systemctl is-active rsyncd.socket rsyncd.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 `rsyncd.socket` and `rsyncd.service`, and remove the `rsync-daemon` package:

```
# systemctl stop rsyncd.socket rsyncd.service
# dnf remove rsync-daemon
```

- OR -

- IF - the `rsync-daemon` package is required as a dependency:

Run the following commands to stop and mask the `rsyncd.socket` and `rsyncd.service`:

```
# systemctl stop rsyncd.socket rsyncd.service
# systemctl mask rsyncd.socket rsyncd.service
```

References:

1. GRID: NIX-00000140
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.		●	●

2.1.14 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 SNMP service is not required, the `net-snmp` 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 `net-snmp` package. If the `net-snmp` package is removed, these packages will be removed as well.

Before removing the `net-snmp` package, review any dependent packages to determine if they are required on the system. If a dependent package is required, stop and mask the `snmpd.service` leaving the `net-snmp` package installed.

Audit:

Run the following command to verify `net-snmp` package is not installed:

```
# rpm -q net-snmp

package net-snmp is not installed
```

- 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 `net-snmp` package:

```
# systemctl stop snmpd.service
# dnf remove net-snmp
```

- 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.		●	●
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.		●	●

2.1.15 Ensure telnet server services are not in use (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `telnet-server` package contains the `telnet` daemon, which accepts connections from users from other systems via the `telnet` protocol.

Rationale:

The `telnet` protocol is insecure and unencrypted. The use of an unencrypted transmission medium could allow a user with access to sniff network traffic the ability to steal credentials. The `ssh` package provides an encrypted session and stronger security.

Impact:

There may be packages that are dependent on the `telnet-server` package. If the `telnet-server` package is removed, these dependent packages will be removed as well. Before removing the `telnet-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 `telnet.socket` leaving the `telnet-server` package installed.

Audit:

Run the following command to verify the `telnet-server` package is not installed:

```
# rpm -q telnet-server

package telnet-server is not installed
```

- OR -

- IF - a package is installed **and** is required for dependencies:

Run the following command to verify `telnet.socket` is not enabled:

```
# systemctl is-enabled telnet.socket 2>/dev/null | grep 'enabled'

Nothing should be returned
```

Run the following command to verify `telnet.socket` is not active:

```
# systemctl is-active telnet.socket 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 `telnet.socket` and remove the `telnet-server` package:

```
# systemctl stop telnet.socket
# dnf remove telnet-server
```

- OR -

- IF - a package is installed **and** is required for dependencies:

Run the following commands to stop and mask `telnet.socket`:

```
# systemctl stop telnet.socket
# systemctl mask telnet.socket
```

References:

1. GRID: NIX-00000143
2. NIST SP 800-53 Rev. 5: CM-7, CM-11
3. NIST SP 800-53A :: CM-7.1 (ii)
4. RHEL 8 STIG Vul ID: V-230487
5. RHEL 8 STIG Rule ID: SV-230487r627750

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.		●	●

2.1.16 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 `tftp-server` package. If the `tftp-server` package is removed, these dependent packages will be removed as well. Before removing the `tftp-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 `tftp.socket` and `tftp.service` leaving the `tftp-server` package installed.

Audit:

Run the following command to verify `tftp-server` is not installed:

```
# rpm -q tftp-server

package tftp-server is not installed
```

- OR -

- IF - the package is required for dependencies:

Run the following command to verify `tftp.socket` and `tftp.service` are not enabled:

```
# systemctl is-enabled tftp.socket tftp.service 2>/dev/null | grep 'enabled'

Nothing should be returned
```

Run the following command to verify the `tftp.socket` and `tftp.service` are not active:

```
# systemctl is-active tftp.socket tftp.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 `tftp.socket` and `tftp.service`, and remove the `tftp-server` package:

```
# systemctl stop tftp.socket tftp.service
# dnf remove tftp-server
```

- OR -

- IF - the `tftp-server` package is required as a dependency:

Run the following commands to stop and mask `tftp.socket` and `tftp.service`:

```
# systemctl stop tftp.socket tftp.service
# systemctl mask tftp.socket tftp.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.		●	●

2.1.17 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.

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` package is not installed:

```
# rpm -q squid

package squid is not installed
```

- 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:

```
# 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
# dnf remove 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-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.		●	●

2.1.18 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 a local site approved requirement to run a web server service on the system, web server packages should be removed to reduce the potential attack surface.

Impact:

Removal of a web server's package, or changing the state of its service and/or socket, will prevent the server from hosting web services.

- IF - a web server package is required for a dependency, any related service or socket should be stopped and masked.

Audit:

Run the following command to verify `httpd` and `nginx` are not installed:

```
# rpm -q httpd nginx

package httpd is not installed
package nginx is not installed
```

- OR -

- IF - a package is installed **and** is required for dependencies:

- IF - the `httpd` package is required for dependencies:

Run the following command to verify `httpd.socket` and `httpd.service` are not enabled:

```
# systemctl is-enabled httpd.socket httpd.service 2>/dev/null | grep
'enabled'

Nothing should be returned
```

Run the following command to verify `httpd.socket` and `httpd.service` are not active:

```
# systemctl is-active httpd.socket httpd.service 2>/dev/null | grep '^active'

Nothing should be returned
```

- IF - the `nginx` package is required for dependencies:

Run the following command to verify `nginx.service` is not enabled:

```
# systemctl is-enabled nginx.service 2>/dev/null | grep 'enabled'

Nothing should be returned
```

Run the following command to verify `nginx.service` is not active:

```
# systemctl is-active nginx.service 2>/dev/null | grep '^active'

Nothing should be returned
```

Note:

- Other web 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 `httpd.socket` and `httpd.service` and remove the `httpd` package:

```
# systemctl stop httpd.socket httpd.service
# dnf remove httpd
```

Run the following commands to stop `nginx.service` and remove the `nginx` package:

```
# systemctl stop nginx.service
# dnf remove nginx
```

- OR -

- IF - a package is installed **and** is required for dependencies:

- IF - the httpd package is required for dependencies:

Run the following commands to stop and mask `httpd.socket` and `httpd.service`:

```
# systemctl stop httpd.socket httpd.service
# systemctl mask httpd.socket httpd.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. If the package is required for a dependency, the service and socket should be stopped and masked.

References:

1. GRID: NIX-00000146
2. NIST SP 800-53 :: 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.		●	●

2.1.19 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:

Configuring the MTA for local-only mode does not affect local mail delivery, but the host will no longer accept mail from other systems. The `postfix` service must be restarted for the change to take effect.

Audit:

This recommendation is designed around postfix, which is the default Mail Transfer Agent (MTA). If an alternative MTA such as exim or sendmail is installed, consult its documentation for the equivalent settings.

If no MTA is installed, the system is compliant with this recommendation.

Verify the MTA is not listening on a non-loopback address:

Run the following command:

```
# ss -plntu
```

Review the output for the standard mail ports 25, 465, and 587. Any of these ports that are in use should be bound only to a loopback address (`127.0.0.1` or `::1`); none should be listening on a non-loopback network interface.

Verify the MTA is bound to the loopback interface only:

Run the following command:

```
# postconf inet_interfaces
```

The returned value should be `loopback-only` (or `localhost`).

For an alternative MTA, use the equivalent command and confirm the MTA is bound only to a loopback address:

- exim: `# exim -bP local_interfaces`
- sendmail: `# grep -i "O DaemonPortOptions=" /etc/mail/sendmail.cf`

Remediation:

Edit `/etc/postfix/main.cf` and add the following line to the RECEIVING MAIL section. If the line already exists, change it to look like the line below:

```
inet_interfaces = loopback-only
```

Run the following command to restart `postfix`:

```
# systemctl restart postfix
```

Note:

- This recommendation is designed around the postfix mail server.
- Depending on your environment you may have an alternative MTA installed such as exim or sendmail. If this is the case consult the documentation for your installed MTA to configure the recommended state.

Default Value:

By default, the `postfix` package sets `inet_interfaces = localhost` in `/etc/postfix/main.cf`, so a default installation listens for mail only on the loopback interface.

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.		●	●

2.1.20 Ensure a graphical user interface is not in use (Automated)

Profile Applicability:

- Level 2 - Server

Description:

A graphical user interface (GUI) is the desktop environment through which users interact with the system using graphical windows and a graphical login. On this platform it is provided by the "Server with GUI" environment group, which installs the display server, a desktop environment, and a display manager such as the GNOME Display Manager (GDM). A system is using a graphical user interface when the "Server with GUI" environment group is installed and the system boots into the graphical target.

Rationale:

If a Graphical User Interface (GUI) is not required, it should be removed to reduce the attack surface of the system.

Impact:

Removing the graphical user interface removes the ability to interact with the system through a local desktop and graphical login. A reboot is required for the change to take effect. Where a graphical interface is a documented operational requirement, this recommendation may not be applicable; before removing it, review any packages that depend on the graphical environment to determine whether they are required on the system.

Audit:

Run the following command to verify the "Server with GUI" environment group is not installed:

```
# dnf grouplist | sed -n '/Installed Environment Groups:/,/Installed
Groups:/p'
```

"Server with GUI" should not be listed under Installed Environment Groups.

- OR -

- IF - the graphical environment is required as a dependency, verify the system does not boot into it — the default target is not the graphical target:

```
# systemctl get-default

multi-user.target
```

Remediation:

Run the following to remove the graphical user interface.

1. Install the "Minimal Install" group to satisfy dependencies (minimizing the impact of removing the graphical group):

```
# dnf groupinstall -y "Minimal Install"
```

2. Remove the "Server with GUI" environment group:

```
# dnf groupremove -y "Server with GUI"
```

3. Install the "Server" group:

```
# dnf groupinstall -y "Server"
```

4. Set the default target to multi-user and reboot:

```
# systemctl set-default multi-user.target
# reboot
```

- OR -

- IF - the graphical environment is required as a dependency but a graphical session is not needed, set the default target so the system does not boot into the graphical interface, then reboot:

```
# systemctl set-default multi-user.target
# reboot
```

Default Value:

Systems installed with the "Server with GUI" environment have a graphical user interface (a display manager and the graphical target) present and in use by default. Systems installed with the "Server" or "Minimal Install" environment have no graphical user interface.

References:

1. GRID: NIX-00000131
2. systemctl(1)

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.		●	●

2.1.21 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-ons. 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:

Removing the X Window server removes the ability to run any local graphical interface; do not apply this on systems where a GUI is required.

If a Graphical Desktop Manager (such as GDM) is in use on the system, there may be a dependency on the `xorg-x11-server-common` package. If the display manager 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 Window `xorg-x11-fonts`. One workaround to avoid this dependency is to use the "headless" Java packages for your specific Java runtime.

Audit:

Note: If a graphical user interface is required on this system (a display manager such as GDM is in use), this recommendation is **Not Applicable**.

- IF - a Graphical Desktop Manager or X Window server is not required and approved by local site policy:

Run the following command to verify the X Window server is not installed:

```
# rpm -q xorg-x11-server-common

package xorg-x11-server-common is not installed
```

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 packages:

```
# dnf remove xorg-x11-server-common
```

References:

1. GRID: NIX-00000147
2. NIST SP 800-53 Rev. 5: CM-11
3. NIST SP 800-53A :: CM-6.1 (iv)
4. RHEL 8 STIG Vul ID: V-230553
5. RHEL 8 STIG Rule ID: SV-230553r809324

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.		●	●

2.1.22 Ensure Xwayland is not in use (Automated)

Profile Applicability:

- Level 2 - Server

Description:

Xwayland is a compatibility X server that allows legacy X11 (Xorg) applications to run within a Wayland session by enabling the X11 protocol for the applications that use it. It is provided by the `xorg-x11-server-Xwayland` package and is started on demand by the Wayland compositor when an X11 client requests a display.

Rationale:

Xwayland reintroduces the X11 protocol. The X11 protocol provides no isolation between clients: any application connected to the X server can read another client's keystrokes, capture the contents of other windows, and inject synthetic input events. Removing Xwayland reduces the X11 attack surface and ensures all graphical applications run under the isolated Wayland protocol.

Impact:

Legacy applications that rely on the X11 protocol will not run once Xwayland is removed. Where `xorg-x11-server-Xwayland` is a required dependency of the graphical environment, `dnf remove` does not remove it in isolation - it cascades to the display manager and the rest of the desktop stack (dozens of packages on a GNOME install), disabling the graphical user interface.

Audit:

1. Run the following command to determine whether the Xwayland package is a required dependency of an installed package:

```
# dnf repoquery --installed --whatrequires xorg-x11-server-Xwayland
```

If any package is returned, `xorg-x11-server-Xwayland` is a required dependency. If a required dependency is listed, and approved by local site policy, it should be noted in the system security plan and this recommendation is not applicable. If nothing is returned, run the following to verify the package is not installed:

2. Run the following command to verify the `xorg-x11-server-Xwayland` package is not installed:

```
# rpm -q xorg-x11-server-Xwayland
```

Example output:

```
package xorg-x11-server-Xwayland is not installed
```

Remediation:

1. Run the following command to confirm removal is limited to the `xorg-x11-server-Xwayland` package:

```
# dnf remove --assumeno xorg-x11-server-Xwayland
```

Only `xorg-x11-server-Xwayland` should be listed. To intentionally remove a graphical environment, remove the environment group (see the graphical user interface recommendation).

2. Run the following command to remove the `xorg-x11-server-Xwayland` package:

```
# dnf remove xorg-x11-server-Xwayland
```

Default Value:

The `xorg-x11-server-Xwayland` package is installed by default on systems that include a graphical desktop environment.

References:

1. GRID: NIX-00000418
2. NIST SP 800-53 :: CM-7
3. NIST SP 800-53 :: CM-11
4. Xwayland(1)

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.		●	●

2.1.23 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
  - **- 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
# dnf remove <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.		●	●

2.2 Configure Client Services

This section covers ensuring insecure network service client packages are not installed unless they are required for an authorized purpose.

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.

Insecure service clients are often used as troubleshooting or testing tools. Removing them can inhibit that capability; where a client is required for a specific task, it is advisable to remove it again after use to prevent accidental or intentional misuse.

Note: This should not be considered a comprehensive list of insecure service clients. You may wish to consider additions to those listed here for your environment.

2.2.1 Ensure ftp client is not installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

FTP (File Transfer Protocol) is a traditional and widely used standard tool for transferring files between a server and clients over a network, especially where no authentication is necessary (permits anonymous users to connect to a server).

Rationale:

FTP does not protect the confidentiality of data or authentication credentials. It is recommended SFTP be used if file transfer is required. Unless there is a need to run the system as a FTP server (for example, to allow anonymous downloads), it is recommended that the package be removed to reduce the potential attack surface.

Impact:

Removing the `ftp` client prevents connecting to FTP servers from this host. Where interactive file transfer is required, use `sftp` or `scp` instead.

Audit:

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

```
# rpm -q ftp

package ftp is not installed
```

Remediation:

Run the following command to remove `ftp`:

```
# dnf remove ftp
```

References:

1. GRID: NIX-00000149
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.		●	●

2.2.2 Ensure ldap client is not installed (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - 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:

Run the following command to verify that the `openldap-clients` package is not installed:

```
# rpm -q openldap-clients

package openldap-clients is not installed
```

Remediation:

Run the following command to remove the `openldap-clients` package:

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

References:

1. GRID: NIX-00000150
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	2.6 Address unapproved software Ensure that unauthorized software is either removed or the inventory is updated in a timely manner	●	●	●

2.2.3 Ensure telnet client is not installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `telnet` package contains the `telnet` client, which allows users to start connections to other systems via the telnet protocol.

Rationale:

The `telnet` protocol is insecure and unencrypted. The use of an unencrypted transmission medium could allow an unauthorized user to steal credentials. The `ssh` package provides an encrypted session and stronger security and is included in most Linux distributions.

Impact:

Removing the `telnet` client prevents establishing outbound telnet connections from this host. Use `ssh` for remote command-line access instead.

Audit:

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

```
# rpm -q telnet

package telnet is not installed
```

Remediation:

Run the following command to remove the `telnet` package:

```
# dnf remove telnet
```

References:

1. GRID: NIX-00000152
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	2.6 Address unapproved software Ensure that unauthorized software is either removed or the inventory is updated in a timely manner	●	●	●

2.2.4 Ensure tftp client is not installed (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:

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:

Removing the `tftp` client prevents TFTP file transfers from this host (for example, fetching network device configuration or firmware images).

Audit:

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

```
# rpm -q tftp

package tftp is not installed
```

Remediation:

Run the following command to remove `tftp`:

```
# dnf remove tftp
```

References:

1. GRID: NIX-00000153
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.		●	●

2.3 Configure Time Synchronization

This section covers ensuring a time synchronization service is in use, configuring it to synchronize from authorized upstream time sources, and running the service securely - for example, not as the `root` user.

It is recommended that systems be configured to synchronize their time using a service such as `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.

Where host-based time synchronization is in use, consult your virtualization platform's documentation to confirm it is active; otherwise, verify that `chrony` is installed and configured to use an authoritative time source.

Only one time synchronization method should be in use on the system. Only the recommendations related to the method in use should be followed; the others should be skipped.

2.3.1 Ensure time synchronization 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: If another method for time synchronization is being used, this section may be skipped.

Rationale:

Time synchronization is important to support time sensitive security mechanisms like Kerberos and also ensures log files have consistent time records across the enterprise, which aids in forensic investigations.

Impact:

Installing `chrony` has negligible operational impact. Accurate, synchronized time is required for reliable log correlation across systems and for time-sensitive security mechanisms such as Kerberos and certificate validation.

Audit:

Run the following command to verify that `chrony` is installed:

```
# rpm -q chrony

chrony-<version>
```

Remediation:

Run the following command to install `chrony`:

```
# dnf install chrony
```

References:

1. GRID: NIX-00000086
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.		●	●

2.3.2 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.

Impact:

Configuring `chrony` with authorized time sources has negligible operational impact and is required for accurate, synchronized time across the enterprise.

Audit:

Run the following command to verify that at least one authorized time source is configured (a `server` or `pool` directive):

```
# grep -Pi -- '^\h*(server|pool)\h+\H+' /etc/chrony.conf
```

A `server` or `pool` line referencing an authorized time source should be returned.

Note: If `/etc/chrony.conf` pulls in additional files with a `confdir` or `sourcedir` directive, also check those included files for the `server`/`pool` directives.

Remediation:

Edit `/etc/chrony.conf` and add or edit `server` and/or `pool` directives to reference authorized time sources according to local site policy. The directive format is:

```
server <remote-server> iburst
pool <remote-pool> iburst
```

For example:

```
server time-a-g.nist.gov iburst
pool 2.pool.ntp.org iburst
```

Run the following command to reload the `chronyd` configuration:

```
# systemctl reload-or-restart chronyd
```

Note: As packaged, `chrony` reads `/etc/chrony.conf` directly and does not enable a general drop-in directory by default (only NetworkManager/DHCP-provided sources via `sourcedir`). Add the `server`/`pool` directives to `/etc/chrony.conf` unless a `confdir` or `sourcedir` has been explicitly configured on the system.

References:

1. GRID: NIX-00000155
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.		●	●

2.3.3 Ensure chrony is not run as the root user (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The file `/etc/sysconfig/chronyd` allows configuration of options for `chrony` to include the user `chrony` is run as. By default, this is set to the user `chrony`

Rationale:

Services should not be set to run as the root user.

Impact:

Running `chronyd` as the dedicated `chrony` user rather than `root` applies the principle of least privilege, limiting the potential impact if the service is compromised.

Audit:

Run the following command to verify that `chrony` isn't configured to run as the `root` user:

```
# grep -Psi -- '^\h*OPTIONS=\"?\h*([^#\n\r]+\h+)?-u\h+root\b'
/etc/sysconfig/chronyd

Nothing should be returned
```

Remediation:

Edit the file `/etc/sysconfig/chronyd` and add or modify the following line to remove "`-u root`" from any `OPTIONS=` argument:

Example:

```
OPTIONS="-F 2"
```

Run the following command to reload the `chronyd.service` configuration:

```
# systemctl reload-or-restart chronyd.service
```

Default Value:

OPTIONS="-F 2"

References:

1. GRID: NIX-00000157

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			

2.4 Job Schedulers

A job scheduler is used to execute jobs, commands, or shell scripts, at fixed times, dates, or intervals

2.4.1 Configure cron

This section covers ensuring the cron daemon is active and that access to the cron configuration files and directories is restricted.

`cron` is a time-based job scheduler.

- IF - `cron` is not installed on the system, this subsection 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.

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.

Impact:

The `cron` daemon must be enabled and active for scheduled jobs to run; ensuring this has no significant operational impact.

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

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			

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.

Impact:

No operational impact. Restricting access to `/etc/crontab` does not affect cron job execution; `crond` runs as root and retains 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

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.	●	●	●

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.

Impact:

No operational impact. Restricting access to `/etc/cron.hourly` does not affect cron job execution; `crond` runs as root and retains 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/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)

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.	●	●	●

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.

Impact:

No operational impact. Restricting access to `/etc/cron.daily` does not affect cron job execution; `crond` runs as root and retains 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/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

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.	●	●	●

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.

Impact:

No operational impact. Restricting access to `/etc/cron.weekly` does not affect cron job execution; `crond` runs as root and retains 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/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)

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.	●	●	●

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.

Impact:

No operational impact. Restricting access to `/etc/cron.monthly` does not affect cron job execution; `crond` runs as root and retains 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/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)

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.	●	●	●

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.

Impact:

No operational impact. Restricting `/etc/cron.yearly` to `root:root 700` does not affect cron job execution; `crond` runs as root and retains 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/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)

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.	●	●	●

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.

Impact:

No operational impact. Restricting access to `/etc/cron.d` does not affect cron job execution; `crond` runs as root and retains 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/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
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.	●	●	●

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.

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:

```
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**.
- 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
```

Default Value:

By default, `/etc/cron.allow` does not exist; `/etc/cron.deny`, if shipped, is empty.

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.	●	●	●

2.4.2 Configure at

This section covers restricting which users may schedule jobs with the `at` command.

`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.

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.

Impact:

No operational impact. Restricting access to the `at` command via `/etc/at.allow` and `/etc/at.deny` does not affect job execution for authorized users.

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`

```
# 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

```
#!/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
}
```

Default Value:

By default, `/etc/at.allow` does not exist; `/etc/at.deny`, if shipped, is empty.

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

3.1 Configure Network Devices

This section covers reducing the attack surface of the system's network hardware - identifying IPv6 status, disabling unused wireless interfaces, and removing or disabling Bluetooth.

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.

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
- Stateful and Stateless configuration
- Support for quality of service (QoS)
- An ideal protocol for neighboring node interaction

Rationale:

IETF RFC 4038 recommends that applications are built with an assumption of dual stack. It is recommended that IPv6 be enabled and configured in accordance with Benchmark recommendations.

- IF - dual stack and IPv6 are not used in your environment, IPv6 may be disabled to reduce the attack surface of the system, and recommendations pertaining to IPv6 can be skipped.

Note: It is recommended that IPv6 be enabled and configured unless this is against local site policy

Impact:

IETF RFC 4038 recommends that applications are built with an assumption of dual stack.

When enabled, IPv6 will require additional configuration to reduce risk to the system.

Audit:

Run the following script to identify if IPv6 is enabled on the system:

```
#!/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:

This recommendation is informational: identify whether IPv6 is enabled and account for it when hardening the system.

- **If IPv6 is enabled** (the expected state on most systems): apply the IPv6-specific network hardening recommendations in this Benchmark.
- **If IPv6 is not required** and no installed software or service depends on it: it may be disabled to reduce the attack surface, in accordance with local site policy. Many modern components rely on IPv6 (including the loopback address `::1`), so test thoroughly before disabling.

Where a change of state is required by local site policy, IPv6 can be disabled by setting the kernel parameters `net.ipv6.conf.all.disable_ipv6 = 1` and `net.ipv6.conf.default.disable_ipv6 = 1` (then reloading with `sysctl --system`), or by adding the `ipv6.disable=1` kernel boot parameter; remove or reverse these to enable IPv6.

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.		●	●

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.

Rationale:

- IF - wireless is not to be used, wireless devices can be disabled to reduce the potential attack surface.

Impact:

Many if not all laptop workstations and some desktop workstations will connect via wireless requiring these interfaces be enabled.

Audit:

Determine whether any wireless interfaces are present:

```
# find /sys/class/net/*/ -type d -name wireless
```

If nothing is returned, the system has no wireless interfaces and is compliant.

If one or more wireless interfaces are present, list their driver modules:

```
# 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 each module returned, verify it is not in use and cannot be loaded:

- `# modprobe -n -v <module>` should return `install /bin/false` (or `/bin/true`) - the module is not loadable
- `# lsmod | grep <module>` should return nothing - the module is not loaded
- `# modprobe --showconfig | grep -P -- '^\h*blacklist\h+<module>\b'` should return a match - the module is deny listed

Remediation:

Run the following script to disable any wireless interfaces:

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

{
   module_fix()
   {
      if ! modprobe -n -v "$l_mname" | grep -P -- '^\h*install
\/bin\/(true|false)'; then
         echo -e " - setting module: \"$l_mname\" to be un-loadable"
         echo -e "install $l_mname /bin/false" >>
/etc/modprobe.d/"$l_mname".conf
      fi
      if lsmod | grep "$l_mname" > /dev/null 2>&1; then
         echo -e " - unloading module \"$l_mname\""
         modprobe -r "$l_mname"
      fi
      if ! grep -Pq -- "^\h*blacklist\h+$l_mname\b" /etc/modprobe.d/*; then
         echo -e " - deny listing \"$l_mname\""
         echo -e "blacklist $l_mname" >> /etc/modprobe.d/"$l_mname".conf
      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_fix
      done
   fi
}
```

Default Value:

By default, wireless interface drivers load automatically when wireless hardware is present.

References:

1. GRID: NIX-00000173
2. NIST SP 800-53 :: AC-18 (1)
3. NIST SP 800-53 :: AC-18 (1)
4. RHEL 8 STIG Vuln ID: V-230506
5. RHEL 8 STIG Rule ID: SV-230506r627750

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.			●

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:

```
# rpm -q bluez

package bluez is not installed
```

- OR -

- IF - the `bluez` package is required as a dependency:

Run the following command to verify `bluetooth.service` is not enabled:

```
# 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
# dnf remove 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

Default Value:

The `bluez` package (which provides `bluetooth.service`) is typically present on desktop/Workstation installations and absent on minimal/Server installations.

References:

1. GRID: NIX-00000171
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.		●	●

3.2 Configure Network Kernel Modules

This section covers ensuring uncommon and legacy network-protocol kernel modules are not available unless explicitly required.

The Linux kernel modules support several network protocols that are not commonly used. If these protocols are not needed, it is recommended that they be disabled in the kernel.

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

3.2.1 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.

Impact:

The Datagram Congestion Control Protocol (DCCP) is rarely used. Disabling the `dccp` kernel module has no operational impact unless an application on the system specifically requires DCCP.

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 module built into the kernel rather than available as a loadable 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:

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

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.		●	●

3.2.2 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.

Impact:

The Transparent Inter-Process Communication (TIPC) protocol is used for communication between nodes in some clusters. Disabling the `tipc` kernel module has no operational impact unless the system participates in a TIPC-based cluster.

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` module built into the kernel rather than available as a loadable 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:

```
# 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
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.		●	●

3.2.3 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.

Impact:

The Reliable Datagram Sockets (RDS) protocol is used by some high-performance workloads (for example, Oracle RAC over InfiniBand). Disabling the `rds` kernel module has no operational impact unless such an application requires 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 module built into the kernel rather than available as a loadable 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:

```
# 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

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.		●	●

3.2.4 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.

Impact:

The Stream Control Transmission Protocol (SCTP) is used by some telephony and signaling stacks (for example SS7 and Diameter). Disabling the `sctp` kernel module has no operational impact unless an application on the system requires SCTP.

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` module built into the kernel rather than available as a loadable 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:

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
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.		●	●

3.2.5 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.

Impact:

Asynchronous Transfer Mode (ATM) is a legacy networking technology. Disabling the `atm` kernel module has no operational impact on modern systems unless ATM hardware or software is specifically in use.

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 module built into the kernel rather than available as a loadable 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:

```
# 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
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.		●	●

3.2.6 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.

Impact:

The Controller Area Network (CAN) protocol is used for automotive and industrial bus communication (SocketCAN). Disabling the `can` kernel module has no operational impact unless the system interfaces with a CAN bus.

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 module built into the kernel rather than available as a loadable 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:

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
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.		●	●

3.3 Configure Network Kernel Parameters

This section covers hardening the kernel parameters that govern the system's network stack, organized into IPv4 and IPv6 subsections. These parameters are managed with `sysctl`.

Kernel parameter (`sysctl`) configuration

- `sysctl` settings are defined in `.conf` files under `/etc/`, `/run/`, `/usr/local/lib/`, `/usr/lib/`, and `/lib/` — typically in a `sysctl.d` subdirectory. Files must use the `.conf` extension. Vendor defaults usually live under `/usr/lib/` or `/usr/local/lib/`.
- `/usr/lib/systemd/systemd-sysctl --cat-config` lists the effective kernel parameters and the file each is configured in. Within that output, a value listed later overrides the same setting listed earlier, so the **last** occurrence of a parameter is the effective (running) value.
- **`--system` precedence.** When run with `--system`, `sysctl` reads files from the directories below in order (top to bottom); once a filename is loaded, a same-named file in a later directory 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. A file may be fully replaced (a same-named file in a higher-priority directory) or partially replaced (a file ordered later).

- **Setting and applying a value.** Place site settings in a `.conf` file under `/etc/sysctl.d/` (the highest-precedence directory, so it reliably overrides vendor defaults), then apply them to the running system with `sysctl --system`. If more than one file sets the same parameter, the effective value is the one that wins by the precedence above. If that effective value is correct but is set in a file **outside** `/etc/sysctl.d/`, add a matching entry under `/etc/sysctl.d/` so a package update cannot silently change it, and comment out or correct any conflicting values in other `/etc/sysctl.d/` files.

3.3.1 Configure IPv4 parameters

This section covers hardening the system's IPv4 network kernel parameters (`net.ipv4.*`), including packet forwarding, ICMP redirect handling, source-route rejection, reverse-path filtering, martian-packet logging, and TCP SYN cookies. These parameters govern how the IPv4 network stack handles routing, redirects, and related security behavior.

3.3.1.1 Ensure net.ipv4.ip_forward is configured (Automated)

Profile Applicability:

- Level 1 - Workstation
- Level 2 - 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.

Note:

- **If this system is a router this recommendation is not applicable.**
- If `net.ipv4.conf.all.forwarding=0` and `net.ipv4.conf.default.forwarding=0` are configured, this recommendation may be skipped.

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`:

```
#!/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)"
   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"
```

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:

```
#!/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)"
   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
```

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

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.		●	●

3.3.1.2 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.

Note:

- **If this system is a router this recommendation is not applicable.**
- If `net.ipv4.ip_forward=0` is configured, this recommendation may be skipped.

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`:

```
#!/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)"
   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"
```

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)"
   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
```

3. Run the following command to load all sysctl configuration files:

```
# sysctl --system
```

Default Value:

net.ipv4.conf.all.forwarding = 1

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.		●	●

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.

Note:

- **If this system is a router this recommendation is not applicable.**
- If `net.ipv4.ip_forward=0` is configured, this recommendation may be skipped.

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:

```
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)"
   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"
```

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)"
   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
```

3. Run the following command to load all sysctl configuration filles:

```
# 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.		●	●

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.

Note: If this system is a router this recommendation is not applicable.

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:

A system acting as a router will no longer send ICMP redirect messages, so directly connected hosts may not be informed of a better first-hop route and could use a suboptimal path. Normal connectivity and packet forwarding are not affected.

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`:

```
#!/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)"
   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"
```

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)"
   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
```

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv4.conf.all.send_redirects = 1

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.		●	●

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.

Note: If this system is a router this recommendation is not applicable.

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:

A system acting as a router will no longer send ICMP redirect messages, so directly connected hosts may not be informed of a better first-hop route and could use a suboptimal path. Normal connectivity and packet forwarding are not affected.

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`:

```
#!/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)"
   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"
```

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)"
   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
```

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv4.conf.default.send_redirects = 1

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.		●	●

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`:

```
#!/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)"
   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"
```

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)"
   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
```

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv4.icmp_ignore_bogus_error_responses = 1

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.		●	●

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.

Impact:

The system will no longer respond to ICMP echo requests sent to broadcast or multicast addresses. Broadcast-based network diagnostics, such as discovering live hosts with a broadcast ping, will not receive a response from this system.

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`:

```
#!/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)"
   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"
```

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)"
   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
```

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv4.icmp_echo_ignore_broadcasts = 1

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.		●	●

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.accept_redirects = 0` is set in a file being used by systemd sysctl to configure `net.ipv4.conf.all.accept_redirects`:

```
#!/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)"
   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"
```

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)"
   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
```

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv4.conf.all.accept_redirects = 1

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.		●	●

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.

Note: If 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.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`:

```
#!/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)"
   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"
```

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)"
   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
```

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv4.conf.default.accept_redirects = 1

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.		●	●

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`:

```
#!/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)"
   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"
```

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)"
   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
```

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv4.conf.all.secure_redirects = 1

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.		●	●

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`:

```
#!/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)"
   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"
```

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)"
   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
```

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv4.conf.default.secure_redirects = 1

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.		●	●

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.

Note: If this system is a router, and uses asymmetrical routing, this recommendation is not applicable.

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:

```
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)"
   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"
```

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)"
   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
```

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv4.conf.all.rp_filter = 2

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.		●	●

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.

Note: If this system is a router, and uses asymmetrical routing, this recommendation is not applicable.

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:

```
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)"
   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"
```

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)"
   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
```

3. Run the following command to load all sysctl configuration filles:

```
# 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.		●	●

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:

```
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)"
   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"
```

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)"
   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
```

3. Run the following command to load all sysctl configuration filles:

```
# 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.		●	●

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:

```
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)"
   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"
```

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)"
   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
```

3. Run the following command to load all sysctl configuration filles:

```
# 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.		●	●

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.

Impact:

Logging packets with impossible ("martian") source or destination addresses increases the volume of kernel log messages. On a busy or heavily spoofed network segment this can consume additional disk space and add noise to the logs.

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`:

```
#!/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)"
   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"
```

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)"
   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
```

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv4.conf.all.log_martians = 0

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.		●	●

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.

Impact:

Logging packets with impossible ("martian") source or destination addresses increases the volume of kernel log messages. On a busy or heavily spoofed network segment this can consume additional disk space and add noise to the logs.

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`:

```
#!/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)"
   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"
```

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)"
   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
```

3. Run the following command to load all sysctl configuration filles:

```
# sysctl --system
```

Default Value:

net.ipv4.conf.default.log_martians = 0

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.		●	●

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.

Impact:

SYN cookies are only engaged when the SYN backlog queue overflows, which typically occurs during a SYN flood. For the connections handled by SYN cookies during such an event, some TCP options negotiated in the initial SYN may not be honored. Connections under normal load are not affected.

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`:

```
#!/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)"
   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"
```

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`.

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)"
   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
```

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
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.		●	●

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.

Impact:

Enforcing `net.ipv4.conf.all.route_localnet = 0` prevents the host from routing or forwarding traffic destined to `127/8`. On systems that legitimately require this — for example a router, or a container/Kubernetes host where `kube-proxy` (or similar) sets `route_localnet=1` — this recommendation may not be applicable; assess and document any such exception before enforcing.

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`:

```
#!/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)"
   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

```
#!/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)"
   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
```

Run the following command to load all sysctl configuration files:

```
# sysctl --system
```

Default Value:

net.ipv4.conf.all.route_localnet = 0

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.		●	●
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.		●	●

3.3.1.20 Ensure net.ipv4.conf.default.route_localnet is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `net.ipv4.conf.default.route_localnet` flag sets, for newly added network interfaces, whether the system may use `127/8` for local routing purposes; when set to `0` the system continues to treat 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.default.route_localnet` to `0` enforces network isolation for newly added interfaces by preventing external interfaces from routing traffic to loopback addresses.

Impact:

Enforcing `net.ipv4.conf.default.route_localnet = 0` prevents newly added interfaces from routing or forwarding traffic destined to `127/8`. On systems that legitimately require this — for example a router, or a container/Kubernetes host where `kube-proxy` (or similar) sets `route_localnet=1` — this recommendation may not be applicable; assess and document any such exception before enforcing.

Audit:

Verify `net.ipv4.conf.default.route_localnet` is set to `0`.

1. Run the following command to verify that `net.ipv4.conf.default.route_localnet` is set to `0` in the running configuration:

```
# sysctl net.ipv4.conf.default.route_localnet
```

Verify output is:

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

2. Run the following script to verify `net.ipv4.conf.default.route_localnet = 0` is set in a file being used by systemd sysctl to configure `net.ipv4.conf.default.route_localnet`:

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

{
   l_parameter_name="net.ipv4.conf.default.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)"
   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.route_localnet = 0`

Example output

```
- "net.ipv4.conf.default.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.default.route_localnet` lines that are not `net.ipv4.conf.default.route_localnet = 0`.

Example script

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

{
   l_option="net.ipv4.conf.default.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)"
   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.default.route_localnet = 0`

Example

```
# [ ! -d "/etc/sysctl.d/" ] && mkdir -p /etc/sysctl.d/
# printf '%s\n' "" "net.ipv4.conf.default.route_localnet = 0" \
>> /etc/sysctl.d/60-ipv4_sysctl.conf
```

Run the following command to load all sysctl configuration files:

```
# sysctl --system
```

Default Value:

net.ipv4.conf.default.route_localnet = 0

3.3.2 Configure IPv6 parameters

This section covers hardening the system's IPv6 network kernel parameters (`net.ipv6.*`), including packet forwarding, ICMP redirect handling, source-route rejection, and router advertisement acceptance. These parameters govern how the IPv6 network stack handles routing, redirects, and related security behavior.

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
```

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:

```
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)"
   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"
```

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)"
   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
```

3. Run the following command to load all sysctl configuration filles:

```
# 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.		●	●

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:

```
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)"
   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"
```

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)"
   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
```

3. Run the following command to load all sysctl configuration filles:

```
# 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.		●	●

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:

```
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)"
   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"
```

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)"
   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
```

3. Run the following command to load all sysctl configuration filles:

```
# 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.		●	●

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:

```
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)"
   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"
```

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)"
   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
```

3. Run the following command to load all sysctl configuration filles:

```
# 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.		●	●

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:

```
# 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)"
   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_source_route = 0`.

Example output:

```
 - "net.ipv6.conf.all.accept_source_route = 0" is set in: "/etc/sysctl.d/60-
ipv6_sysctl.conf"
```

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)"
   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
```

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.		●	●

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:

```
# 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)"
   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"
```

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)"
   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:

```
# [ ! -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
```

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.		●	●

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.

Impact:

On networks that rely on IPv6 Router Advertisements (SLAAC) for autoconfiguration, no longer accepting Router Advertisements will prevent the system from automatically configuring its IPv6 address, default route, and gateway. Such systems must instead use static IPv6 addressing or DHCPv6.

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`:

```
#!/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)"
   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"
```

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)"
   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
```

3. Run the following command to load all sysctl configuration filles:

```
# 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.		●	●

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.

Impact:

On networks that rely on IPv6 Router Advertisements (SLAAC) for autoconfiguration, no longer accepting Router Advertisements will prevent the system from automatically configuring its IPv6 address, default route, and gateway. Such systems must instead use static IPv6 addressing or DHCPv6.

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`:

```
#!/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)"
   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"
```

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)"
   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
```

3. Run the following command to load all sysctl configuration filles:

```
# 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 Host Based Firewall, on a Linux system, is a set of rules. When a data packet moves into or out of a protected network space, its contents (in particular, information about its origin, target, and the protocol it plans to use) are tested against the firewall rules to see if it should be allowed through.

To provide a Host Based Firewall, the Linux kernel includes support for nftables.

nftables is a 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.

In order to configure firewall rules for nftables, a firewall utility needs to be installed. `firewalld` is the default and most commonly used host-based firewall on the platforms this Benchmark targets, and is the utility for which guidance is provided in this section. It provides firewall features by acting as a front-end for the Linux kernel's netfilter framework via the nftables backend, as an alternative to the `nft` command-line program. `firewalld` supports both IPv4 and IPv6 networks and can administer separate firewall zones with varying degrees of trust as defined in zone profiles.

This Benchmark includes guidance for configuring nftables through the `firewalld` 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 firewalld's ability to translate IPTables commands to nftables.
- This section is intended only to ensure the resulting firewall rules are in place.
- The `ipset` and `iptables-nft` packages are deprecated. This includes deprecation of nft-variants such as `iptables`, `ip6tables`, `arptables`, and `ebtables` utilities. If you are using any of these tools, for example, because you upgraded from an earlier version, we recommend migrating to using the nftables package for backend services.

4.1 Configure firewalld

This section covers configuring the `firewalld` host-based firewall — installing it and enabling its service, selecting its backend, setting zone targets and a default-drop zone, protecting loopback traffic, permitting only required services and ports, logging denied packets, blocking ICMP timestamp requests, and ensuring no conflicting firewall utilities are in use.

`firewalld` uses the concepts of **zones** and **services** to simplify traffic management. A zone defines the trust level applied to the interfaces and source addresses assigned to it and controls which incoming traffic is allowed. A service is a named set of ports and protocols (for example `ssh`) that can be permitted within a zone.

Important: Configuration of a live system's firewall directly over a remote connection will often result in being locked out. It is advised to have a known good firewall configuration set to run on boot and to configure an entire firewall structure in a script that is then run and tested before saving to boot.

The following example will create a `firewalld` zone called `securezone` to implement the firewall rules of this section leveraging the firewalld utility included with the firewalld package. This example will open port 22 (ssh) from anywhere. Opening service `SSH` should be updated in accordance with local site policy. If another name for the zone is preferred, replace `securezone` with the name to be used.

Sample securezone zone xml file

```
<?xml version="1.0" encoding="utf-8"?>
<zone target="DROP">
  <description>For use with CIS Linux Benchmark. You do not trust the other
computers on networks to not harm your computer. Only selected incoming
connections are accepted.</description>
  <service name="ssh"/>
  <service name="dhcpv6-client"/>
  <icmp-block name="destination-unreachable"/>
  <icmp-block name="packet-too-big"/>
  <icmp-block name="time-exceeded"/>
  <icmp-block name="parameter-problem"/>
  <icmp-block name="neighbor-advertisement"/>
  <icmp-block name="neighbor-solicitation"/>
  <icmp-block name="router-advertisement"/>
  <icmp-block name="router-solicitation"/>
  <rule family="ipv4">
    <source address="127.0.0.0/8"/>
    <destination address="127.0.0.0/8" invert="True"/>
    <drop/>
  </rule>
  <rule family="ipv6">
    <source address="::1"/>
    <destination address="::1" invert="True"/>
    <drop/>
  </rule>
  <icmp-block-inversion/>
</zone>
```

Note: To use this zone, save this as `/etc/firewalld/zones/securezone.xml` and run the following commands:

```
# firewall-cmd --reload
# firewall-cmd --permanent --zone=securezone --change-interface={NAME OF
NETWORK INTERFACE}
```

4.1.1 Ensure firewalld is installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`firewalld` is a high-level firewall rules management frontend utility. It does not directly handle packet filtering itself but relies on the `nftables` subsystem within the Linux kernel.

Rationale:

`firewalld` is a dynamic, user friendly, firewall manager that can protect against threats originating from within a corporate network including malicious mobile code and poorly configured software on a host.

Impact:

Installing the `firewalld` package alone does not change firewall behavior; rules take effect only after the service is enabled and a zone/policy is configured.

Audit:

Run the following command to verify that `firewalld` is installed:

```
# rpm -q firewalld

firewalld-<version>
```

Remediation:

Run the following command to install `firewalld`

```
# dnf install firewalld
```

References:

1. GRID: NIX-00000174
2. NIST SP 800-53 Rev. 5: CA-9

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.4 Implement and Manage a Firewall on Servers Implement and manage a firewall on servers, where supported. Example implementations include a virtual firewall, operating system firewall, or a third-party firewall agent.	●	●	●
v7	9.4 Apply Host-based Firewalls or Port Filtering Apply host-based firewalls or port filtering tools on end systems, with a default-deny rule that drops all traffic except those services and ports that are explicitly allowed.	●	●	●

4.1.2 Ensure firewalld backend is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `FirewallBackend` option selects the firewall backend implementation.

Choices are:

- nftables (default)
- iptables (iptables, ip6tables, ebtables and ipset)

Rationale:

The `iptables` backend is deprecated in favor of `nftables`, the modern default packet-filtering framework that `firewalld` uses.

Impact:

On a system already using the default `nftables` backend, no change occurs. Changing the backend from `iptables` requires a `firewalld` reload and rebuilds the active ruleset using nftables.

Audit:

`nftables` is the default `firewalld` backend, so no explicit configuration is required. The backend is only non-compliant when it is explicitly set to `iptables`.

Run the following command to verify the backend is not set to `iptables`:

```
# grep -Pi -- '^\h*FirewallBackend\h*=\h*iptables\b'
/etc/firewalld/firewalld.conf
```

Nothing should be returned.

Remediation:

Edit the file `/etc/firewalld/firewalld.conf` and add or modify the following line:

```
FirewallBackend=nftables
```

Default Value:

FirewallBackend=nftables

References:

1. GRID: NIX-00000175
2. NIST SP 800-53 Rev. 5: CA-9

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.4 Implement and Manage a Firewall on Servers Implement and manage a firewall on servers, where supported. Example implementations include a virtual firewall, operating system firewall, or a third-party firewall agent.	●	●	●
v7	9.4 Apply Host-based Firewalls or Port Filtering Apply host-based firewalls or port filtering tools on end systems, with a default-deny rule that drops all traffic except those services and ports that are explicitly allowed.	●	●	●

4.1.3 Ensure firewalld.service is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`firewalld.service` is a firewall service daemon that provides a dynamic, customizable firewall with a D-Bus interface.

Rationale:

`firewalld.service` must be active to enforce rules configured through FirewallD. `firewalld.service` must be enabled to start automatically after a system reboot.

Impact:

Enabling and starting `firewalld.service` activates host-based packet filtering: from that point, incoming traffic is evaluated against the applicable zone and any traffic that is not permitted is blocked. Ensure the services and ports required for the system's role — including remote administration such as SSH — are allowed, or connectivity to the system can be affected.

Audit:

Run the following command to verify `firewalld.service` is enabled:

```
# systemctl is-enabled firewalld.service

enabled
```

Run the following command to verify `firewalld.service` is active:

```
# systemctl is-active firewalld.service

active
```

Remediation:

Run the following commands to unmask, enable, and start `firewalld.service`:

```
# systemctl unmask firewalld.service
# systemctl --now enable firewalld.service
```

Default Value:

`firewalld.service` is enabled by default (systemd preset) on systems where the `firewalld` package is installed.

References:

1. GRID: NIX-00000176

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.4 Implement and Manage a Firewall on Servers Implement and manage a firewall on servers, where supported. Example implementations include a virtual firewall, operating system firewall, or a third-party firewall agent.	●	●	●
v7	9.4 Apply Host-based Firewalls or Port Filtering Apply host-based firewalls or port filtering tools on end systems, with a default-deny rule that drops all traffic except those services and ports that are explicitly allowed.	●	●	●

4.1.4 Ensure firewalld active zone target is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

FirewallD zones are a network or firewall zone that defines the trust level of the interface used for a connection. There are several pre-defined zones provided by firewalld.

A zone target may be set to `ACCEPT`, `REJECT`, or `DROP`. This can be used to accept, reject, or drop every packet that doesn't match any rule (port, service, etc). The `ACCEPT` target is used in the `trusted` zone to accept every packet not matching any rule, e.g. accept all. The `REJECT` target is used in the `block` zone to reject, with default firewall reject type, every packet not matching any rule. The `DROP` target is used in the `drop` zone to drop every packet not matching any rule. If the target is not specified, every packet not matching any rule will be rejected.

Note: The target may also be set to `default`. The `default` target provides a convenient way to implement a "deny by default, allow by exception" policy within a firewalld zone, while still allowing essential ICMP traffic.

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:

Setting an active zone's target to `DROP` or `REJECT` causes traffic that does not match an explicit allow rule to be dropped or rejected. Ensure the services and ports required for the system's role — including remote administration such as SSH — are permitted, or connectivity to the system can be lost.

Audit:

Run the following command to verify the active zone `target` is not set to `ACCEPT`:

Run the following script to verify:

- active zone `target` for interfaces that are not loopback or virtual is not set to `ACCEPT`.
- an active zone is configured.
- active zone target is written as a permanent firewall rule.

```
{
   a_output=() a_output2=()
   while IFS= read -r l_zone; do
      l_zone="$(awk -F'(' '{print $1}' <<< "${l_zone// /}")"
      l_interfaces="$(firewall-cmd --zone="$l_zone" --list-interfaces)"
      if ! grep -Psq -- '^(lo|virbr[^\h]+)\h*$' <<< "${l_interfaces/ /}";
then
         l_ptarget="$(firewall-cmd --permanent --zone="$l_zone" --get-
target)"
         l_target="$(cut -d: -f2 <<< "$(grep -Pso -- '^\h*target:\h+\H+\b' \
         <<< "$(firewall-cmd --list-all --zone="$l_zone")")" | xargs)"
         if [ -z "$l_target" ] || [ "${l_target,,}" = "accept" ]; then
            a_output2+=(" - Active zone: \"$l_zone\"" \
            "    Target is: \"$l_target\" for interfaces: \"$l_interfaces\"")
         elif [ ! "${l_target,,}" = "${l_ptarget,,}" ]; then
            a_output2+=(" - Active zone: \"$l_zone\"" \
            "   Target is: \"$l_target\" is not a permanent firewall rule")
         else
            a_output+=(" - Active zone: \"$l_zone\"" \
            "    Target is: \"$l_target\""
            "    for interfaces: \"$l_interfaces\"")
         fi
      fi
   done < <(firewall-cmd --get-active-zones | awk '$1!~/interface/{print}')
   printf '%s\n' "" "${a_output2[@]}" "" "${a_output[@]}" ""
}
```

Review the script output. The audit passes when all of the following are true:

- at least one active zone (excluding the loopback and virtual interfaces) is reported (empty output means no active zone is configured and does not meet this Recommendation);
- no reported zone has a `target` of `ACCEPT`; and
- each reported zone's `target` is also written as a permanent firewall rule.

The script lists any non-compliant zones first (a `target` of `ACCEPT`, or a `target` that is not written as a permanent rule); no such entries should be present.

Remediation:

Run the following command to update an active zone target and write a permanent firewalld rule:

```
# firewall-cmd --zone={ACTIVE_ZONE} --set-target={TARGET} --permanent
```

Example:

```
# firewall-cmd --zone=public --set-target=DROP --permanent
```

Note: `--set-target=default` is also acceptable.

References:

1. GRID: NIX-00000177

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.	●	●	●

4.1.5 Ensure firewalld loopback traffic is configured (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Every device with networking capabilities has a loopback interface.

Loopback traffic refers to network communication where a device sends data to itself, essentially routing the traffic back to its own network interface.

Rationale:

Loopback traffic is generated between processes on machine and is typically critical to operation of the system. The loopback interface is the only place that loopback network traffic should be seen, all other interfaces should ignore traffic on this network as an anti-spoofing measure.

Audit:

Verify the loopback interface accepts all incoming traffic.

1. Run the following command to verify the currently configured zone for the loopback interface:

```
# firewall-cmd --get-zone-of-interface=lo
```

Example output:

```
trusted
```

Verify the output is not `no zone`. If the output is `no zone` the second audit is also automatically considered to be failed.

2. Run the following command to verify the configured zone for the loopback interface `target` is set to `ACCEPT`:

```
# firewall-cmd --list-all --zone="$(firewall-cmd --get-zone-of-interface=lo |
xargs)" | \
grep -Psi -- '^\h*target:\h+\H+'
```

Verify output is:

```
  target: ACCEPT
```

Note: The commands above query the runtime `firewalld` configuration. Ensure the setting is also written as a permanent rule (the Remediation uses `--permanent`); a runtime-only change will not persist across a `firewall-cmd --reload` or a reboot.

Remediation:

1. Run one of the following commands to add the `lo` interface to the trusted zone.

If the `lo` interface is not currently assigned to a zone:

```
# firewall-cmd --permanent --zone=trusted --add-interface=lo
```

- OR - if the `lo` interface is currently assigned to a zone:

```
# firewall-cmd --permanent --zone=trusted --change-interface=lo
```

2. Run the following command to reload firewalld:

```
# firewall-cmd --reload
```

References:

1. GRID: NIX-00000178

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.	●	●	●

4.1.6 Ensure firewalld loopback source address traffic is configured (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Every device with networking capabilities has a loopback interface.

Loopback traffic refers to network communication where a device sends data to itself, essentially routing the traffic back to its own network interface.

Rationale:

The loopback interface is the only place that loopback network traffic should be seen, all other interfaces should ignore traffic on this network as an anti-spoofing measure.

Audit:

Run the following script to verify active zones include dropping traffic with a loopback address as the source address and not also the destination address:

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

{
   while IFS= read -r l_zone; do
      echo " - Zone: $l_zone"
      firewall-cmd --zone=$l_zone --list-rich-rules | \
      grep -P -- '(::1|127\.0\.0\.0/8)\b'
   done < <(firewall-cmd --get-zone-of-interface=lo | xargs)
}
```

Verify the output includes:

- A zone.
- The rule: `rule family="ipv4" source address="127.0.0.0/8" destination NOT address="127.0.0.0/8" drop`
- The rule: `rule family="ipv6" source address="::1" destination NOT address="::1" drop`

Example output:

```
 - Zone: trusted
rule family="ipv4" source address="127.0.0.0/8" destination NOT
address="127.0.0.0/8" drop
rule family="ipv6" source address="::1" destination NOT address="::1" drop
```

Note: The command above queries the runtime `firewalld` configuration. Ensure the rich rules are also written as permanent rules (the Remediation uses `--permanent`); runtime-only rules will not persist across a `firewall-cmd --reload` or a reboot.

Remediation:

Add the appropriate rich rules by running the following commands:

```
# firewall-cmd --permanent --zone={ZONE_NAME} --add-rich-rule='rule
family=ipv4 source address="127.0.0.0/8" destination not
address="127.0.0.0/8" drop'

# firewall-cmd --permanent --zone={ZONE_NAME} --add-rich-rule='rule
family=ipv6 source address="::1" destination not address="::1" drop'
```

Example script:

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

{
   while IFS= read -r l_zone; do
      firewall-cmd --permanent --zone="$l_zone" --add-rich-rule='rule
family=ipv4 source address="127.0.0.0/8" destination not
address="127.0.0.0/8" drop'
      firewall-cmd --permanent --zone="$l_zone" --add-rich-rule='rule
family=ipv6 source address="::1" destination not address="::1" drop'
   done < <(firewall-cmd --get-zone-of-interface=lo | xargs)
}
```

Run the following command to load the updated firewalld rules:

```
# firewall-cmd --reload
```

References:

1. GRID: NIX-00000179

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.	●	●	●

4.1.7 Ensure firewalld services and ports are configured (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`firewalld` controls which traffic a zone allows by permitting named **services** (predefined protocol/port groups such as `ssh` or `dhcpv6-client`) and individual **ports**. Traffic that is not matched by an allowed service or port is handled by the zone's target (accept, reject, or drop). Only the services and ports required for the system's role should be permitted, keeping the exposed attack surface to a minimum.

Rationale:

To reduce the attack surface of a system, all services and ports should be blocked unless required

Impact:

Removing a service or port that is required for the system's role will break access to it. Confirm each permitted service and port against site policy before removing, and validate that remote access (such as SSH) is retained.

Audit:

Run the following command and review output to ensure that listed services and ports follow site policy:

```
# firewall-cmd --list-all --zone="$(firewall-cmd --list-all | awk
'/\(active\)/ { print $1 }')" | grep -P -- '^\h*(services:|ports:)'
```

Note: The command `firewall-cmd --get-services | tr ' ' '\n' | column` can be used to produce a list of services firewalld can be configured to address as services.

Remediation:

Run the following command to remove an unnecessary service:

```
# firewall-cmd --remove-service=<service>
```

Example:

```
# firewall-cmd --remove-service=cockpit
```

Run the following command to remove an unnecessary port:

```
# firewall-cmd --remove-port=<port-number>/<port-type>
```

Example:

```
# firewall-cmd --remove-port=25/tcp
```

Run the following command to make new settings persistent:

```
# firewall-cmd --runtime-to-permanent
```

References:

1. GRID: NIX-00000180
2. https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html-single/configuring_firewalls_and_packet_filters/index

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.4 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.	●	●	●

4.1.8 Ensure firewalld default zone is set to drop (Manual)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

`firewalld` should be set to deny all traffic by default.

Rationale:

A deny-all default reduces the host attack surface. With the default zone set to `drop`, incoming traffic on any interface not assigned to another zone is dropped unless it is explicitly permitted, so only intended services are reachable.

Impact:

Setting the default zone to `drop` causes incoming traffic on interfaces not assigned to another zone to be dropped unless explicitly allowed; required services (including remote administration such as SSH) must be permitted or connectivity can be lost. Outbound traffic is not affected by the default zone.

Note: Configuration of a live system's 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.

Audit:

Run the following command to verify the default zone is set to `drop`:

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

drop
```

`drop` should be returned.

Remediation:

Run the following to set the drop zone as the default.

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

Reload firewalld for changes to take effect

```
# firewall-cmd --reload
```

Default Value:

By default, the `firewalld` default zone is `public`.

References:

1. GRID: NIX-00000181

4.1.9 Ensure conflicting firewall utilities are not in use with firewalld (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Running more than one host-based firewall manager on a system can produce conflicting or unpredictable rule sets. `firewalld` manages the Linux kernel's netfilter framework through the NFTables backend; a separately managed `nftables` ruleset (via `nftables.service`), the legacy `iptables-services` saved rules, or `ufw` running alongside `firewalld` can override or conflict with the firewalld-managed rules. Only a single host-based firewall utility should be active on the system.

Rationale:

When multiple firewall management utilities are active, the effective rule set becomes ambiguous and difficult to audit, and one utility can silently override the protective rules configured in another - weakening or negating the host firewall. Ensuring `firewalld` is the only active firewall manager keeps the enforced policy deterministic and auditable.

Impact:

If a conflicting utility is currently providing required firewall rules, disabling or removing it without first migrating those rules into `firewalld` could remove needed access controls. Migrate any required rules into firewalld before disabling other firewall managers.

Audit:

Run the following checks to verify that no conflicting firewall manager is in use alongside `firewalld`.

1. Verify the standalone `nftables.service` is not enabled or active (firewalld programs the NFTables backend directly and does not use this service):

```
# systemctl is-enabled nftables.service
# systemctl is-active nftables.service
```

Neither command should return `enabled` or `active` (`disabled`, `masked`, `inactive`, or a not-found result is expected).

2. Verify the legacy `iptables-services` package is not installed:

```
# rpm -q iptables-services
```

`package iptables-services is not installed` should be returned.

3. Verify `ufw` is not installed:

```
# rpm -q ufw
```

`package ufw is not installed` should be returned.

Remediation:

For any conflicting firewall manager found in the Audit, migrate any required rules into `firewalld` first, then disable or remove it.

Stop, disable, and mask a standalone `nftables.service`:

```
# systemctl stop nftables.service
# systemctl --now disable nftables.service
# systemctl mask nftables.service
```

Remove the `iptables-services` package if installed:

```
# dnf remove iptables-services
```

Remove the `ufw` package if installed:

```
# dnf remove ufw
```

Reload firewalld to apply the firewalld-managed ruleset:

```
# firewall-cmd --reload
```

Default Value:

`firewalld` uses the NFTables backend directly; `nftables.service` is not enabled by default, and `iptables-services` and `ufw` are not installed by default.

4.1.10 Ensure firewalld logs denied packets (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

The `LogDenied` option in `/etc/firewalld/firewalld.conf` controls logging of packets denied (dropped or rejected) by `firewalld`. Supported values are `off` (default), `all`, `unicast`, `broadcast`, and `multicast`. When enabled, firewalld adds kernel logging rules for denied packets immediately before the applicable reject/drop rules.

Rationale:

Logging denied packets provides visibility into blocked and attempted connections, supporting detection of firewall misconfigurations and malicious activity, and providing evidence for incident investigation.

Impact:

Enabling `LogDenied` increases the volume of kernel/firewall log messages; ensure log rotation and storage are sized appropriately, as high-traffic or heavily-probed hosts can generate a large number of entries.

Audit:

Run the following command to verify `firewalld` is logging denied packets:

```
# firewall-cmd --get-log-denied
```

The value returned should be one of `all`, `unicast`, `broadcast`, or `multicast` (any value other than `off`), in accordance with local site policy.

Remediation:

Run the following command to enable logging of denied packets (replace `all` with the value that meets local site policy):

```
# firewall-cmd --set-log-denied=all
```

Run the following command to reload firewalld:

```
# firewall-cmd --reload
```

Default Value:

`LogDenied=off`

4.1.11 Ensure firewalld blocks ICMP timestamp requests and replies (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

ICMP timestamp messages (`timestamp-request`, ICMP type 13, and `timestamp-reply`, ICMP type 14) allow a remote host to query the system clock. `firewalld` can block these message types, per zone, so the host does not respond to timestamp requests.

Rationale:

Answering ICMP timestamp requests discloses the host's system clock, which can help an attacker fingerprint the host, correlate activity across hosts, or attack time-dependent protocols. Vulnerability scanners and several STIGs flag hosts that respond to ICMP timestamp requests. Blocking these ICMP types removes the disclosure.

Impact:

ICMP timestamp messages are rarely used by legitimate applications, so blocking them typically has no operational effect. Diagnostic tools that rely on ICMP timestamp (uncommon) will no longer function.

Audit:

ICMP timestamp blocking applies to `firewalld` zones that filter network traffic. A zone whose target is `ACCEPT` (such as the built-in `trusted` zone) does not filter traffic and is not applicable; a zone whose target is `DROP` or `%%REJECT%%` already discards these messages. Only a zone whose target is `default` needs the explicit block.

1. Run the following command to list the active zones:

```
# firewall-cmd --get-active-zones
```

2. Run the following command to determine each active zone's target:

```
# firewall-cmd --permanent --zone=<ZONE> --get-target
```

3. For each active zone whose target is `default`, check the ICMP block inversion setting and list the zone's ICMP blocks:

```
# firewall-cmd --permanent --zone=<ZONE> --query-icmp-block-inversion
# firewall-cmd --permanent --zone=<ZONE> --list-icmp-blocks
```

If inversion is not enabled (`no`), the icmp-block list is a block-list: `timestamp-request` and `timestamp-reply` should be listed. If inversion is enabled (`yes`), the icmp-block list is an allow-list: `timestamp-request` and `timestamp-reply` should not be listed (any type not listed is blocked).

Example (target `default`, inversion disabled):

```
# firewall-cmd --permanent --zone=public --get-target
default
# firewall-cmd --permanent --zone=public --list-icmp-blocks
timestamp-reply timestamp-request
```

Remediation:

Block the ICMP timestamp message types on each active zone that filters network traffic (target `default`). A zone whose target is `ACCEPT` (such as `trusted`) or `DROP`/`%%REJECT%%` already handles these messages and needs no change.

1. For each such zone, block both message types:

```
# firewall-cmd --permanent --zone=<ZONE> --add-icmp-block=timestamp-request
# firewall-cmd --permanent --zone=<ZONE> --add-icmp-block=timestamp-reply
```

2. Reload firewalld to apply the change:

```
# firewall-cmd --reload
```

Example (blocking on the `public` zone):

```
# firewall-cmd --permanent --zone=public --add-icmp-block=timestamp-request
# firewall-cmd --permanent --zone=public --add-icmp-block=timestamp-reply
# firewall-cmd --reload
```

Default Value:

By default, `firewalld` does not block ICMP timestamp messages; the host answers `timestamp-request` messages.

5 Access Control

This section provides guidance for controlling access to the system and for authenticating and authorizing users. It covers hardening the OpenSSH server, configuring privilege escalation, the Pluggable Authentication Modules (PAM) stack, and user account and environment defaults. Review the overview of each subsection for setting-specific guidance before applying its recommendations.

5.1 Configure SSH Server

This section provides recommendations for hardening the OpenSSH server (sshd) — protecting access to its configuration and host key files, ensuring strong cryptographic algorithms, and setting the authentication, session, and logging controls that reduce the risk of unauthorized remote access.

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. To audit a setting that may be governed by a `Match` block, re-run the audit supplying the relevant connection parameters, for example: `# sshd -T -C user=<user> | grep <keyword>`.
- 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 `Ciphers`, `MACs`, or `KexAlgorithms` arguments
    - 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 an 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 `Ciphers`, `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.
  - 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, RDomain, 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.

5.1.1 Ensure access to /etc/ssh/sshd_config is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The file `/etc/ssh/sshd_config`, and files ending in `.conf` in the `/etc/ssh/sshd_config.d` directory, contain configuration specifications for `sshd`.

Rationale:

Configuration specifications for `sshd` need to be protected from unauthorized access by non-privileged users.

Impact:

Restricting `/etc/ssh/sshd_config` and its drop-in `.conf` files to mode `0600`, owner and group `root`, removes read access for non-root users and processes. Any configuration-management, monitoring, or user workflow that reads the SSH server configuration without elevated privileges will need to run with root privileges.

Audit:

Run the following commands to verify that `/etc/ssh/sshd_config` and the `.conf` files in the `/etc/ssh/sshd_config.d` directory are mode `0600` or more restrictive and are owned by the user and group `root`.

Verify the main configuration file. Nothing should be returned:

```
# find /etc/ssh/sshd_config -maxdepth 0 \( -perm /177 -o ! -user root -o ! -
group root \) 2>/dev/null
```

Verify the drop-in `.conf` files. Nothing should be returned:

```
# find /etc/ssh/sshd_config.d/ -type f -name '*.conf' \( -perm /177 -o ! -
user root -o ! -group root \) 2>/dev/null
```

Note: Mode `0600` or more restrictive means the file grants no group or other access and no execute permission for the owner. If other locations are listed in an `Include` statement, the `.conf` files in those 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)
}
```

- 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

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.	●	●	●

5.1.2 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 the recommendation that governs the parent /etc/ssh/sshd_config file and the .conf drop-in files inside this directory, which does not constrain the directory entry itself.

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:

- Enumerate the per-host SSH server policy
- Read the contents of any .conf drop-in that itself is world-readable
- 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.

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. SSHD_CONFIG(5)
4. SSHD(8)

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 unauthorized user obtains the private SSH host key file, the host could be impersonated.

Impact:

sshd enforces the ownership and permissions of host private keys, so an incorrect combination can cause a key to be ignored, and if no usable host key remains, sshd will not start. Acceptable access is owner root with either group root and mode 0600 or more restrictive, or (where host-based authentication via `ssh-keysign` is used) group ssh_keys and mode 0640 or more restrictive; over-restricting a host-based-authentication key to `root:root 0600` will disable that authentication.

Audit:

Verify that each SSH private host key file is owned by user root and is either:

- group-owned by root with mode 0600 or more restrictive, or
- group-owned by ssh_keys with mode 0640 or more restrictive (the ssh_keys group is used only where the distribution provides it, for host-based authentication).

Run the following script to list any private host key file that does not meet the above. Host key files are read from the running sshd configuration (so keys in non-default locations or with non-default names are included), and mode is tested with a permission bitmask (so more-restrictive modes also pass). Nothing should be returned:

```
#!/usr/bin/env bash
{
   l_ssh_group=""
   grep -Psiq -- '^\h*ID_LIKE="?([^#\n\r]+)?\bfedora\b' /etc/os-release
2>/dev/null && \
      grep -Psq -- '^ssh_keys\b' /etc/group 2>/dev/null &&
l_ssh_group="ssh_keys"
   while IFS= read -r l_file; do
      ssh-keygen -lf "$l_file" &>/dev/null || continue
      while IFS=: read -r l_mode l_owner l_group; do
         [ "$l_group" = "$l_ssh_group" ] && l_mask="0137" || l_mask="0177"
         if [ $(( l_mode & l_mask )) -gt 0 ] || [ "$l_owner" != "root" ] || \
            { [ "$l_group" != "$l_ssh_group" ] && [ "$l_group" != "root" ];
}; then
            printf '%s\n' "\"$l_file\": mode $l_mode owner $l_owner group
$l_group"
         fi
      done < <(stat -Lc '%#a:%U:%G' "$l_file")
   done < <(sshd -T 2>/dev/null | awk '$1=="hostkey"{print $2}')
}
```

Remediation:

Run the following script to set the owner and restrict permissions on each SSH private host key file from the running sshd configuration:

```
#!/usr/bin/env bash
{
   while IFS= read -r l_file; do
      ssh-keygen -lf "$l_file" &>/dev/null || continue
      chown root:root "$l_file"
      chmod u-x,go-rwx "$l_file"
   done < <(sshd -T 2>/dev/null | awk '$1=="hostkey"{print $2}')
}
```

- OR - Where host-based authentication requires the ssh_keys group to read the host keys, set group ownership to ssh_keys with mode 0640 or more restrictive instead:

```
#!/usr/bin/env bash
{
   while IFS= read -r l_file; do
      ssh-keygen -lf "$l_file" &>/dev/null || continue
      chown root:ssh_keys "$l_file"
      chmod u-x,g-wx,o-rwx "$l_file"
   done < <(sshd -T 2>/dev/null | awk '$1=="hostkey"{print $2}')
}
```

Default Value:

Access: (0600/-rw-------) Uid: (0/root) Gid: (0/root)

Note: On platforms that provide the `ssh_keys` group, private host keys may instead be group owned by `ssh_keys` with mode `0640` (-rw-r-----).

References:

1. GRID: NIX-00000189
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. CCI-000366
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. STIG ID: RHEL-08-010490 | RULE ID: SV-230287r1017098 | 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.	●	●	●
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.	●	●	●

5.1.4 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.

Impact:

Public host key files are intended to be readable; enforcing mode 0644, owner and group root, only removes excess write and execute permissions and does not affect normal SSH operation.

Audit:

Verify that each SSH public host key file is owned by user root and group root, with mode 0644 or more restrictive.

Run the following script to list any public host key file that does not meet the above. Public host key files are the ".pub" of each host key in the running sshd configuration, and mode is tested with a permission bitmask (so more-restrictive modes also pass). Nothing should be returned:

```
#!/usr/bin/env bash
{
   l_ssh_group=""
   grep -Psiq -- '^\h*ID_LIKE="?([^#\n\r]+)?\bfedora\b' /etc/os-release
2>/dev/null && \
      grep -Psq -- '^ssh_keys\b' /etc/group 2>/dev/null &&
l_ssh_group="ssh_keys"
   while IFS= read -r l_key; do
      l_file="${l_key}.pub"
      ssh-keygen -lf "$l_file" &>/dev/null || continue
      while IFS=: read -r l_mode l_owner l_group; do
         if [ $(( l_mode & 0133 )) -gt 0 ] || [ "$l_owner" != "root" ] || \
            { [ "$l_group" != "$l_ssh_group" ] && [ "$l_group" != "root" ];
}; then
            printf '%s\n' "\"$l_file\": mode $l_mode owner $l_owner group
$l_group"
         fi
      done < <(stat -Lc '%#a:%U:%G' "$l_file")
   done < <(sshd -T 2>/dev/null | awk '$1=="hostkey"{print $2}')
}
```

Remediation:

Run the following script to set the owner, group, and restrict permissions on each SSH public host key file from the running sshd configuration:

```
#!/usr/bin/env bash
{
   while IFS= read -r l_key; do
      l_file="${l_key}.pub"
      [ -e "$l_file" ] || continue
      chown root:root "$l_file"
      chmod u-x,go-wx "$l_file"
   done < <(sshd -T 2>/dev/null | awk '$1=="hostkey"{print $2}')
}
```

Default Value:

644 0/root 0/root

References:

1. GRID: NIX-00000190
2. NIST SP 800-53A :: CM-6.1 (iv)
3. NIST SP 800-53 Rev. 5: AC-3 CM-6 b MP-2
4. RHEL 8 STIG Rule ID: SV-230286r627750

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.	●	●	●

5.1.5 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 script to list any weak ciphers that are enabled. Nothing should be returned:

```
#!/usr/bin/env bash
{
   a_weak=(3des-cbc blowfish-cbc cast128-cbc aes128-cbc aes192-cbc aes256-cbc
arcfour arcfour128 arcfour256 rijndael-cbc@lysator.liu.se)
   l_offered="$(sshd -T | awk '$1=="ciphers"{print $2}' | tr ',' '\n')"
   for l_alg in "${a_weak[@]}"; do
      grep -Fxq -- "$l_alg" <<< "$l_offered" && printf '%s\n' "$l_alg"
   done
}
```

Note: (advisory - chacha20-poly1305 / CVE-2023-48795): The following command is an advisory check only and is not part of the pass/fail assessment above. It reports whether the chacha20-poly1305@openssh.com cipher is offered:

```
# sshd -T 2>&1 | awk '($1=="ciphers" && $2~/chacha20-
poly1305@openssh\.com/){print}'
```

On OpenSSH 9.6 and later the Terrapin attack (CVE-2023-48795) is mitigated by default through strict key exchange, so a returned line does not by itself indicate a weakness. On earlier OpenSSH releases, if a line is returned, review CVE-2023-48795 and verify the system has been patched, and consider excluding chacha20-poly1305 through the system-wide crypto policy.

Remediation:

The ciphers available to the SSH server are governed by the system-wide crypto policy. Disable the weak ciphers by configuring the system-wide crypto policy as described in the "Configure system wide crypto policy" section.

- IF - CVE-2023-48795 (Terrapin) has not been reviewed and addressed, chacha20-poly1305 should also be disabled through the system-wide crypto policy.

Once the crypto policy has been updated, reload the SSH daemon:

```
# systemctl reload-or-restart sshd
```

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. SSHD_CONFIG(5)
7. NIST SP 800-53 Rev. 5: SC-8
8. https://nvd.nist.gov/vuln/detail/CVE-2023-48795

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.		●	●

5.1.6 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

Notes:

- Kex algorithms have a higher preference the earlier they appear in the list
- Some organizations may have stricter requirements for approved Key exchange algorithms
- Ensure that Key exchange algorithms used are in compliance with site policy
- The only Key Exchange Algorithms currently FIPS 140 approved are:
  - ecdh-sha2-nistp256
  - ecdh-sha2-nistp384
  - ecdh-sha2-nistp521
  - diffie-hellman-group-exchange-sha256
  - diffie-hellman-group16-sha512
  - diffie-hellman-group18-sha512
  - diffie-hellman-group14-sha256

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

Impact:

A too restrictive list of key exchange algorithms can lead to a "no matching key exchange method found" error during connection establishment, when a client and the server share no common algorithm. Verify that peers required to connect to this host support at least one of the retained algorithms before applying.

Audit:

Run the following script to list any weak key exchange algorithms that are enabled. Nothing should be returned:

```
#!/usr/bin/env bash
{
   a_weak=(diffie-hellman-group1-sha1 diffie-hellman-group14-sha1 diffie-
hellman-group-exchange-sha1)
   l_offered="$(sshd -T | awk '$1=="kexalgorithms"{print $2}' | tr ',' '\n')"
   for l_alg in "${a_weak[@]}"; do
      grep -Fxq -- "$l_alg" <<< "$l_offered" && printf '%s\n' "$l_alg"
   done
}
```

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:

The key exchange algorithms available to the SSH server are governed by the system-wide crypto policy. The weak key exchange algorithms in scope are SHA1-based; disable them by disabling SHA1 hash and signature support in the system-wide crypto policy as described in the "Configure system wide crypto policy" section.

Once the crypto policy has been updated, reload the SSH daemon:

```
# systemctl reload-or-restart sshd
```

Default Value:

KexAlgorithms sntrup761x25519-sha512@openssh.com,curve25519-sha256,curve25519-sha256@libssh.org,ecdh-sha2-nistp256,ecdh-sha2-nistp384,ecdh-sha2-nistp521,diffie-hellman-group-exchange-sha256,diffie-hellman-group16-sha512,diffie-hellman-group18-sha512,diffie-hellman-group14-sha256

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 key exchange algorithms supported by a given host depend on the installed OpenSSH version and the active system-wide crypto policy. Current OpenSSH releases on this platform support algorithms such as:

curve25519-sha256 curve25519-sha256@libssh.org ecdh-sha2-nistp256 ecdh-sha2-nistp384 ecdh-sha2-nistp521 diffie-hellman-group-exchange-sha256 diffie-hellman-group16-sha512 diffie-hellman-group18-sha512 diffie-hellman-group14-sha256 sntrup761x25519-sha512@openssh.com

Newer OpenSSH releases additionally support the post-quantum hybrid key exchange mlkem768x25519-sha256. Run "sshd -T | grep -i kexalgorithms" to view the algorithms active on the host.

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.		●	●

5.1.7 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.

Notes:

- Some organizations may have stricter requirements for approved MACs.
- Ensure that MACs used are in compliance with site policy.
- The only "strong" MACs currently FIPS 140 approved are:
  - HMAC-SHA1
  - HMAC-SHA2-256
  - HMAC-SHA2-384
  - HMAC-SHA2-512

Rationale:

MD5-based, RIPEMD-160, 96-bit-truncated, and 64-bit (UMAC-64) 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:

A too restrictive list of MAC algorithms can lead to a "no matching MAC found" error during connection establishment, when a client and the server share no common algorithm. Verify that peers required to connect to this host support at least one of the retained MACs before applying.

Audit:

Run the following script to list any weak MACs that are enabled. Nothing should be returned:

```
#!/usr/bin/env bash
{
   a_weak=(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)
   l_offered="$(sshd -T | awk '$1=="macs"{print $2}' | tr ',' '\n')"
   for l_alg in "${a_weak[@]}"; do
      grep -Fxq -- "$l_alg" <<< "$l_offered" && printf '%s\n' "$l_alg"
   done
}
```

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:

```
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
```

Remediation:

The MAC algorithms available to the SSH server are governed by the system-wide crypto policy. Disable the weak MACs by configuring the system-wide crypto policy as described in the "Configure system wide crypto policy" section.

- IF - CVE-2023-48795 (Terrapin) has not been reviewed and addressed, the Encrypt-then-MAC (EtM) MACs should also be disabled through the system-wide crypto policy.

Once the crypto policy has been updated, reload the SSH daemon:

```
# systemctl reload-or-restart 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 :: CM-6, SC-8, SC-8(1), SC-13
5. https://nvd.nist.gov/vuln/detail/CVE-2023-48795

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.		●	●

5.1.8 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.

Rationale:

Restricting which users can remotely access the system via SSH will help ensure that only authorized users access the system.

Impact:

Restricting SSH access to specific users or groups means accounts that are not on the allow list (or that appear on the deny list) can no longer log in over SSH. Ensure the allow/deny lists include every account that legitimately requires remote access before applying, or authorized administrators may be locked out.

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

Note: This setting can be overridden within a Match block. If Match blocks are configured, also verify the effective value for the relevant contexts as described in the section overview.

Remediation:

Edit the `/etc/ssh/sshd_config` file to set one or more of the parameters above any `Include` and `Match` set statements as follows:

```
AllowUsers <userlist>
 - AND/OR -
AllowGroups <grouplist>
```

Note:

- **Be advised** that 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.
- 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.

Default Value:

None

References:

1. GRID: NIX-00000191
2. NIST SP 800-53 Rev. 5: AC-3. MP-2
3. SSHD(8)
4. 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.	●	●	●

5.1.9 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.

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.

Impact:

Configuring a banner presents a pre-authentication notice to every user before login. Automated SSH clients, scripts, or monitoring that do not expect banner text should be verified to still operate correctly, and the banner content must be maintained in line with organizational policy.

Audit:

Run the following command to verify `Banner` is set:

```
# sshd -T | grep -Pi -- '^banner\h+\/\H+'
```

Example:

```
banner /etc/issue.net
```

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}')"
```

Note: This setting can be overridden within a Match block. If Match blocks are configured, also verify the effective value for the relevant contexts as described in the section overview.

Remediation:

Edit the `/etc/ssh/sshd_config` file to 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}')"
```

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			

5.1.10 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.

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.

Impact:

With ClientAliveInterval and ClientAliveCountMax configured, connections that stop responding at the protocol level are disconnected after the configured interval. Long-running but idle interactive sessions, or sessions over unreliable links, may be terminated and need to be re-established; set the values in line with site policy.

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
```

Note: This setting can be overridden within a Match block. If Match blocks are configured, also verify the effective value for the relevant contexts as described in the section overview.

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the `ClientAliveInterval` and `ClientAliveCountMax` parameters above any `Include` and `Match` entries according to site policy.

Example:

```
ClientAliveInterval 15
ClientAliveCountMax 3
```

Default Value:

ClientAliveInterval 0

ClientAliveCountMax 3

References:

1. GRID: NIX-00000194
2. SSHD_CONFIG(5)
3. SSHD(8)
4. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

Additional Information:

https://bugzilla.redhat.com/show_bug.cgi?id=1873547

https://github.com/openssh/openssh-portable/blob/V_8_9/serverloop.c#L137

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			

5.1.11 Ensure sshd DisableForwarding is enabled (Automated)

Profile Applicability:

- Level 1 - Workstation
- Level 2 - 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.

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 intermediate server can make free use of ssh-agent to authenticate as the user 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.

Impact:

Enabling `DisableForwarding` turns off all SSH forwarding features (X11, ssh-agent, TCP, and StreamLocal). Environments that rely on SSH tunneling — for example, to add encryption to legacy applications, to traverse firewalls, or to run remote X11 applications — will lose that capability and must adopt an alternative before this recommendation is applied.

Audit:

Run the following command to verify `DisableForwarding` is set to `yes`:

```
# sshd -T | grep -i disableforwarding

disableforwarding yes
```

Note: This setting can be overridden within a Match block. If Match blocks are configured, also verify the effective value for the relevant contexts as described in the section overview.

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the `DisableForwarding` parameter to `yes` above any `Include` entry as follows:

```
DisableForwarding yes
```

References:

1. GRID: NIX-00000196
2. SSHD(8)
3. sshd_config(5)
4. 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.		●	●

5.1.12 Ensure sshd GSSAPIAuthentication is disabled (Automated)

Profile Applicability:

- Level 1 - Workstation
- Level 2 - Server

Description:

The `GSSAPIAuthentication` parameter specifies whether user authentication based on GSSAPI is allowed

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

Audit:

Run the following command to verify `GSSAPIAuthentication` is set to `no`:

```
# sshd -T | grep gssapiauthentication

gssapiauthentication no
```

Note: This setting can be overridden within a Match block. If Match blocks are configured, also verify the effective value for the relevant contexts as described in the section overview.

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the `GSSAPIAuthentication` parameter to `no` above any `Include` and `Match` entries as follows:

```
GSSAPIAuthentication no
```

Default Value:

GSSAPIAuthentication no

References:

1. GRID: NIX-00000197
2. SSHD(8)
3. SSHD_CONFIG(5)
4. NIST SP 800-53A :: CM-6.1 (iv)
5. RHEL 8 STIG Vul ID: V-244528
6. RHEL 8 STIG Rule ID: SV-244528r952106
7. NIST SP 800-53 :: CM-6 b

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.		●	●

5.1.13 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.

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
```

Note: This setting can be overridden within a Match block. If Match blocks are configured, also verify the effective value for the relevant contexts as described in the section overview.

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the `HostbasedAuthentication` parameter to `no` above any `Include` and `Match` entries as follows:

```
HostbasedAuthentication no
```

Default Value:

HostbasedAuthentication no

References:

1. GRID: NIX-00000198
2. SSHD(8)
3. SSHD_CONFIG(5)
4. NIST SP 800-53A :: CM-6.1 (iv)
5. NIST SP 800-53 :: CM-6 b

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.		●	●

5.1.14 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`.

Rationale:

Setting this parameter forces users to enter a password when authenticating with SSH.

Impact:

Enabling IgnoreRhosts prevents .rhosts and .shosts files from being used during authentication. Any workflow that relies on rhosts-style host-based trust will stop working and must be migrated to key-based authentication before this recommendation is applied.

Audit:

Run the following command to verify `IgnoreRhosts` is set to `yes`:

```
# sshd -T | grep ignorerhosts

ignorerhosts yes
```

Note: This setting can be overridden within a Match block. If Match blocks are configured, also verify the effective value for the relevant contexts as described in the section overview.

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the `IgnoreRhosts` parameter to `yes` above any `Include` entry as follows:

```
IgnoreRhosts yes
```

Default Value:

IgnoreRhosts yes

References:

1. GRID: NIX-00000199
2. SSHD(8)
3. SSHD_CONFIG(5)
4. NIST SP 800-53 Rev. 5: CM-1,CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.2 Use Unique Passwords Use unique passwords for all enterprise assets. Best practice implementation includes, at a minimum, an 8-character password for accounts using MFA and a 14-character password for accounts not using MFA.	●	●	●
v7	4.4 Use Unique Passwords Where multi-factor authentication is not supported (such as local administrator, root, or service accounts), accounts will use passwords that are unique to that system.		●	●

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 section the Grace Period should be limited to appropriate organizational limits to ensure the service is available for needed access.

Rationale:

Setting the `LoginGraceTime` parameter to a low number will minimize the risk of successful brute force attacks to the SSH server. It will also limit the number of concurrent unauthenticated connections. While the recommended setting is 60 seconds (1 minute), set the number based on site policy.

Impact:

Reducing `LoginGraceTime` shortens the window in which a connecting client must complete authentication. On slow or high-latency links, or where interactive multi-factor prompts take time to complete, legitimate users may occasionally exceed a very low grace time and have to reconnect; set the value in line with 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:

Edit the `/etc/ssh/sshd_config` file to set the `LoginGraceTime` parameter to `60` seconds or less above any `Include` entry as follows:

```
LoginGraceTime 60
```

Default Value:

LoginGraceTime 120

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			

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.

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.

Impact:

Setting LogLevel to INFO or VERBOSE increases the volume of SSH log data; VERBOSE additionally records the key fingerprint used for each login. Ensure log storage, rotation, and any downstream log collection are sized to accommodate the additional volume.

Audit:

Run the following command and verify that output matches `loglevel VERBOSE` or `loglevel INFO`:

```
# sshd -T | grep loglevel

loglevel VERBOSE
   - OR -
loglevel INFO
```

Note: This setting can be overridden within a Match block. If Match blocks are configured, also verify the effective value for the relevant contexts as described in the section overview.

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the `LogLevel` parameter to `VERBOSE` or `INFO` above any `Include` and `Match` entries as follows:

```
LogLevel VERBOSE
   - OR -
LogLevel INFO
```

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.		●	●

5.1.17 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.

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.

Impact:

Lowering MaxAuthTries reduces the number of authentication attempts permitted per connection. Users whose clients try several keys before the correct one, or who mistype a passphrase, may reach the limit sooner and have to reconnect; set the value in line with site policy.

Audit:

Run the following command and verify that `MaxAuthTries` is `4` or less:

```
# sshd -T | grep maxauthtries

maxauthtries 4
```

Note: This setting can be overridden within a Match block. If Match blocks are configured, also verify the effective value for the relevant contexts as described in the section overview.

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the `MaxAuthTries` parameter to `4` or less above any `Include` and `Match` entries as follows:

```
MaxAuthTries 4
```

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.			●

5.1.18 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.

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.

Impact:

Tightening MaxStartups causes the daemon to begin randomly dropping new unauthenticated connections earlier as concurrent connection attempts rise. On hosts that legitimately receive many simultaneous connections, tune the value so that authorized clients are not dropped under normal load.

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:

Edit the `/etc/ssh/sshd_config` file to set the `MaxStartups` parameter to `10:30:60` or more restrictive above any `Include` entries as follows:

```
MaxStartups 10:30:60
```

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)

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			

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.

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.

Impact:

Limiting MaxSessions caps the number of open sessions (channels) permitted per SSH connection. Users or tools that multiplex many sessions over a single connection may reach the limit and need to open additional connections; set the value in line with site policy.

Audit:

Run the following command and verify that `MaxSessions` is `10` or less:

```
# sshd -T | grep -i 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
```

Note: This setting can be overridden within a Match block. If Match blocks are configured, also verify the effective value for the relevant contexts as described in the section overview.

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the `MaxSessions` parameter to `10` or less above any `Include` and `Match` entries as follows:

```
MaxSessions 10
```

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			

5.1.20 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.

Rationale:

Disallowing remote shell access to accounts that have an empty password reduces the probability of unauthorized access to the system.

Impact:

This setting is the default. Ensuring PermitEmptyPasswords is no prevents login to accounts that have an empty password; any account intentionally configured without a password will be unable to authenticate over SSH and must be given a credential or an alternative access method.

Audit:

Run the following command to verify `PermitEmptyPasswords` is set to `no`:

```
# sshd -T | grep permitemptypasswords

permitemptypasswords no
```

Note: This setting can be overridden within a Match block. If Match blocks are configured, also verify the effective value for the relevant contexts as described in the section overview.

Remediation:

Edit `/etc/ssh/sshd_config` and set the `PermitEmptyPasswords` parameter to `no` above any `Include` and `Match` entries as follows:

```
PermitEmptyPasswords no
```

The SSH daemon must be restarted for the changes to take effect. To restart the SSH daemon, run the following command:

```
# systemctl reload-or-restart sshd.service
```

Default Value:

PermitEmptyPasswords no

References:

1. GRID: NIX-00000207
2. NIST SP 800-53 Revision 5 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. RHEL 8 STIG Vul ID: V-230380
5. RHEL 8 STIG Rule ID: SV-230380r951612
6. 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.		●	●

5.1.21 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`.

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.

Impact:

With `PermitRootLogin` set to `no`, the root account cannot authenticate directly over SSH. Administrators must connect with an individual account and then escalate privileges. Any automation or process that logs in directly as root over SSH will need to be reworked to use a named account before this recommendation is applied.

Audit:

Run the following command to verify `PermitRootLogin` is set to `no`:

```
# sshd -T | grep permitrootlogin

permitrootlogin no
```

Note: This setting can be overridden within a Match block. If Match blocks are configured, also verify the effective value for the relevant contexts as described in the section overview.

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the `PermitRootLogin` parameter to `no` above any `Include` and `Match` entries as follows:

```
PermitRootLogin no
```

Default Value:

PermitRootLogin prohibit-password

References:

1. GRID: NIX-00000208
2. NIST SP 800-53 Rev. 5:AC-6
3. NIST SP 800-53A :: IA-2 (5).2 (ii)
4. RHEL 8 STIG Vul ID: V-230296
5. RHEL 8 STIG Rule ID: SV-230296r951608
6. 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.	●	●	●

5.1.22 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.

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)

Impact:

Disabling PermitUserEnvironment prevents users from setting environment options at login (for example through a ~/.ssh/environment file). Workflows that rely on user-supplied SSH environment variables will need to use an alternative mechanism before this recommendation is applied.

Audit:

Run the following command to verify `PermitUserEnvironment` is set to `no`:

```
# sshd -T | grep permituserenvironment

permituserenvironment no
```

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the `PermitUserEnvironment` parameter to `no` above any `Include` entries as follows:

```
PermitUserEnvironment no
```

Default Value:

PermitUserEnvironment no

References:

1. GRID: NIX-00000209
2. SSHD(8)
3. NIST SP 800-53 Rev. 5: CM-1,CM-2, CM-6, CM-7, IA-5
4. NIST SP 800-53A :: CM-6.1 (iv)
5. RHEL 8 STIG Vul ID: V-230330
6. RHEL 8 STIG Rule ID: SV-230330r951610
7. 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			

5.1.23 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.

Rationale:

When `UsePAM` is set to `yes`, PAM runs through account and session types properly. This is important if you want to restrict access to services based off of IP, time or other factors of the account. Additionally, you can make sure users inherit certain environment variables on login or disallow access to the server

Impact:

Enabling `UsePAM` (the platform default) has no adverse operational impact on a standard configuration. `UsePAM` must remain enabled for PAM-based account and session processing — including access restrictions, resource limits, and session setup — to take effect; disabling it is not supported on this platform.

Audit:

Run the following command to verify `UsePAM` is set to `yes`:

```
# sshd -T | grep -i usepam

usepam yes
```

Remediation:

Edit the `/etc/ssh/sshd_config` file to set the `UsePAM` parameter to `yes` above any `Include` entries as follows:

```
UsePAM yes
```

Default Value:

UsePAM yes

Note: The upstream OpenSSH compiled default for `UsePAM` is `no`; on this platform `UsePAM` is set to `yes` by `/etc/ssh/sshd_config.d/50-redhat.conf`. Setting `UsePAM no` is not supported on this platform.

References:

1. GRID: NIX-00000210
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.		●	●

5.2 Configure privilege escalation

This section provides recommendations for configuring privilege escalation using sudo. The recommendations ensure sudo is installed, restrict the ownership and permissions of the sudoers configuration files and the /etc/sudoers.d drop-in directory, harden sudo Defaults options (such as password handling and credential caching), and restrict use of the su command.

sudo reads its configuration from /etc/sudoers, which normally ends with an `#includedir /etc/sudoers.d` directive that pulls in the files under that directory; the effective configuration is the combination of /etc/sudoers and the files in /etc/sudoers.d. When the same `Defaults` setting appears more than once, the last matching definition takes effect, so a value set in a later-parsed drop-in file can override an earlier one. When auditing or remediating any setting below, account for all of these files (not just /etc/sudoers), and always edit sudoers files with `visudo`, which validates syntax before saving so a mistake cannot lock out privilege escalation.

5.2.1 Ensure sudo is installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`sudo` allows a permitted user to execute a command as the superuser or another user, as specified by the security policy. The invoking user's real (not effective) user ID is used to determine the user name with which to query the security policy.

Rationale:

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

The security policy determines what privileges, if any, a user has to run `sudo`. The policy may require that users authenticate themselves with a password or another authentication mechanism. If authentication is required, `sudo` will exit if the user's password is not entered within a configurable time limit. This limit is policy-specific.

Audit:

Verify that the sudo package is installed. Run the following command:

```
# rpm -q sudo
```

The output should list the installed package, for example:

```
sudo-<version>
```

Remediation:

Run the following command to install sudo

```
# dnf install sudo
```

References:

1. GRID: NIX-00000215
2. NIST SP 800-53 Rev. 5: AC-6(2), AC-6(5)
3. SUDO(8)

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.	●	●	●

5.2.2 Ensure access on sudoers configuration files is configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The sudo configuration is defined in /etc/sudoers and in any files included from /etc/sudoers.d/. These files determine which users may escalate privileges and under what conditions.

Rationale:

If a sudoers file is writable by a user other than root, that user could grant unauthorized privilege escalation; if it is readable beyond root, the escalation policy is disclosed. Restricting each sudoers file to mode 0440 or more restrictive, owned by root and group root, ensures only root can read or modify the escalation policy. sudo refuses to run if these files are group- or world-writable.

Audit:

Verify that /etc/sudoers and every file in /etc/sudoers.d/ is mode 0440 or more restrictive and is owned by user root and group root. Run the following command to list any file that does not meet these requirements:

```
# find /etc/sudoers /etc/sudoers.d/ -type f \( -perm /0337 -o ! -user root -o
! -group root \)
```

Nothing should be returned.

Remediation:

Set the owner, group, and permissions on the sudoers files so that only root can read or modify them:

```
# chown root:root /etc/sudoers
# chmod u-wx,g-wx,o-rwx /etc/sudoers
# find /etc/sudoers.d/ -type f -exec chown root:root {} +
# find /etc/sudoers.d/ -type f -exec chmod u-wx,g-wx,o-rwx {} +
```

Modify sudoers content only with `visudo`.

Default Value:

/etc/sudoers and files in /etc/sudoers.d/ are mode 0440, owner root, group root, by default.

References:

1. GRID: NIX-00000423
2. SUDOERS(5)
3. SUDO(8)

CIS Controls:

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

5.2.3 Ensure access on the /etc/sudoers.d directory is configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

/etc/sudoers.d/ is the drop-in directory sudo includes via the `#includedir` directive in /etc/sudoers. Files placed there become part of the effective sudo policy.

Rationale:

If /etc/sudoers.d/ is writable by a user other than root, that user could add a drop-in file granting unauthorized privilege escalation; if it is world-accessible, the presence and names of its policy files are exposed. Restricting the directory to mode 0750 or more restrictive, owned by root and group root, ensures only root can add, remove, or list drop-in policy files.

Audit:

Verify that the /etc/sudoers.d/ directory is mode 0750 or more restrictive and is owned by user root and group root. Run the following command:

```
# find /etc/sudoers.d -maxdepth 0 -type d \( -perm /0027 -o ! -user root -o ! -group root \)
```

Nothing should be returned.

Remediation:

Set the owner, group, and permissions on the directory so that only root can add, remove, or list drop-in policy files:

```
# chown root:root /etc/sudoers.d
# chmod g-w,o-rwx /etc/sudoers.d
```

Default Value:

/etc/sudoers.d/ is mode 0750, owner root, group root, by default.

References:

1. GRID: NIX-00000424
2. SUDOERS(5)
3. SUDO(8)

CIS Controls:

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

5.2.4 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`).

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.

Audit:

Verify that `sudo` can only run other commands from a pseudo terminal.

Run the following command to verify `Defaults use_pty` is set:

```
# grep -rPi -- '^\h*Defaults\h+([^#\n\r]+,\h*)?use_pty\b' /etc/sudoers
/etc/sudoers.d/
```

Verify the output matches:

```
/etc/sudoers:Defaults use_pty
```

Run the following command to verify `Defaults !use_pty` is not set:

```
# grep -rPi -- '^\h*Defaults\h+([^#\n\r]+,\h*)?!use_pty\b' /etc/sudoers
/etc/sudoers.d/
```

Nothing should be returned.

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`

Because sudo applies the last matching `Defaults` definition, place this setting so that no later-parsed file in `/etc/sudoers.d` overrides it; see this section's overview for how the sudoers files combine.

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. VISUDO(8)
4. sudoers(5)
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.	●	●	●
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.	●	●	●

5.2.5 Ensure sudo log file exists (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `Defaults logfile` entry sets the path to the sudo log file. Setting a path turns on logging to a file; negating this option turns it off. By default, sudo logs via syslog.

Rationale:

Defining a dedicated log file for sudo simplifies auditing of sudo commands and creation of auditd rules for sudo.

Impact:

Creation of additional log files can cause disk space exhaustion if not correctly managed. You should configure `logrotate` to manage the sudo log in accordance with your local policy.

Audit:

Verify that sudo is configured to write to a custom log file. Run the following command:

```
# grep -rPsi
"^\h*Defaults\h+([^#]+,\h*)?logfile\h*=\h*(\"|\')?\H+(\"|\')?(,\h*\H+\h*)*\h*
(#.*)?$" /etc/sudoers /etc/sudoers.d/
```

The output should show a `logfile` entry, for example:

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

Remediation:

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

```
Defaults  logfile="<PATH TO CUSTOM LOG FILE>"
```

Example:

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

Because sudo applies the last matching `Defaults` definition, place this setting so that no later-parsed file in `/etc/sudoers.d` overrides it; see this section's overview for how the sudoers files combine.

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.
- 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.		●	●

5.2.6 Ensure users must provide password for escalation (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

sudo can be configured to require users to authenticate before running commands with elevated privileges. The `NOPASSWD` tag, where present in the sudoers configuration, lets a user run sudo commands without entering a password; without it, sudo prompts for the user's password before escalating privileges.

Rationale:

The NOPASSWD tag allows a user to run commands via sudo without re-entering a password. Requiring a password for privilege escalation verifies the invoking user's identity before elevated commands run, and limits an attacker with access to an unlocked session or a compromised account from escalating privileges without knowing the account password.

Impact:

Applying this removes passwordless sudo, so service or automation accounts that rely on non-interactive (passwordless) privilege escalation (for example, database failover or orchestration tooling) will no longer be able to escalate without providing a password.

Audit:

Verify that users must supply a password for privilege escalation - that is, no sudoers entry uses the `NOPASSWD` tag.

Note: If passwords are not being used for authentication, this is not applicable.

Run the following command to check `/etc/sudoers` and the files in `/etc/sudoers.d/`:

```
# grep -r "^[^#].*NOPASSWD" /etc/sudoers /etc/sudoers.d/
```

Nothing should be returned. A returned line indicates a `NOPASSWD` tag is in use; refer to the remediation procedure.

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.

Because sudo applies the last matching `Defaults` (and per-rule) definition, ensure no later-parsed file in `/etc/sudoers.d` re-introduces a `NOPASSWD` tag; see this section's overview for how the sudoers files combine.

References:

1. GRID: NIX-00000218
2. NIST SP 800-53 Rev. 5: AC-6

Additional Information:

Some service or automation accounts legitimately require passwordless sudo. Where operationally necessary, grant it narrowly: limited to the specific command(s) and the specific account (and runas target where possible) rather than a blanket NOPASSWD, and treated as a documented, risk-accepted exception with compensating controls (the account restricted to non-interactive/service use, sudo activity logged and monitored, and host/network isolation). The audit still flags the entry, so the exception remains visible and owned.

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.	●	●	●

5.2.7 Ensure re-authentication for privilege escalation is not disabled globally (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

By default sudo re-authenticates the invoking user before escalating privileges. The `!authenticate` tag, where present in the sudoers configuration, disables that re-authentication; without it, sudo prompts the user to confirm their identity before running elevated commands.

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:

Where !authenticate was in use, affected users will be prompted for their password when escalating privileges rather than escalating without one. This likewise affects any non-interactive or automation context that relied on !authenticate to escalate without a prompt.

Audit:

Verify that re-authentication for privilege escalation is not disabled - that is, no sudoers entry uses the `!authenticate` tag.

Run the following command to check `/etc/sudoers` and the files in `/etc/sudoers.d/`:

```
# grep -r '^[^#].*!authenticate' /etc/sudoers /etc/sudoers.d/
```

Nothing should be returned. A returned line with a `!authenticate` tag indicates re-authentication has been disabled; refer to the remediation procedure.

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).

Because sudo applies the last matching definition, ensure no later-parsed file in `/etc/sudoers.d` re-introduces an `!authenticate` tag; see this section's overview for how the sudoers files combine.

References:

1. GRID: NIX-00000213
2. NIST SP 800-53 Rev. 5: AC-6

Additional Information:

Some service or automation accounts legitimately require non-interactive privilege escalation. Where operationally necessary, grant it narrowly: limited to the specific command(s) and the specific account rather than disabling re-authentication broadly, and treated as a documented, risk-accepted exception with compensating controls (the account restricted to non-interactive/service use, sudo activity logged and monitored, and host/network isolation). The audit still flags the entry, so the exception remains visible and owned.

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.	●	●	●

5.2.8 Ensure sudo pwfeedback is configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `pwfeedback` option provides visual feedback (for example, an asterisk) for each character as the sudo password is typed. When not enabled (the default), no feedback is displayed.

Rationale:

Visual feedback discloses the exact length of the password to anyone able to observe the terminal or a session recording, weakening it against shoulder-surfing and offline analysis. The option has also been associated with input-handling defects in older versions of sudo. Leaving `pwfeedback` disabled removes the length side-channel and keeps sudo on its hardened default.

Audit:

Verify that the sudo pwfeedback option is not enabled (this is the secure default). Run the following command to review the sudoers configuration:

```
# grep -rPi -- '^\h*Defaults\h+([^#]+,\h*)?pwfeedback\b' /etc/sudoers
/etc/sudoers.d/
```

Nothing should be returned.

Remediation:

Ensure sudo does not display feedback while the password is typed. Edit the appropriate sudoers file with `visudo` and remove any line that enables pwfeedback, or explicitly disable it:

```
# visudo
```

Add or correct the line:

```
Defaults !pwfeedback
```

Default Value:

pwfeedback is not enabled (disabled) by default.

References:

1. GRID: NIX-00000421
2. SUDOERS(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.	●	●	●

5.2.9 Ensure sudo visiblepw is configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

The `visiblepw` sudo option controls whether sudo will prompt for a password on a terminal that cannot disable local echo. When not enabled (the default), sudo refuses to run rather than prompt for a password that would be displayed as it is typed.

Rationale:

If `visiblepw` is enabled, sudo prompts for a password even when the terminal cannot suppress echo (for example, a command run over SSH without a pseudo-terminal). The password is then entered in cleartext where it can be observed, captured in a screen or session recording, or written to logs. Leaving `visiblepw` disabled ensures a password is only requested when it can be entered without being displayed.

Impact:

Automated or non-interactive workflows that invoke sudo without a pseudo-terminal will fail with a password-prompt error rather than exposing the password. Allocate a pseudo-terminal (for example, `ssh -t`) for such workflows, or use an appropriate documented exception.

Audit:

Verify that the sudo visiblepw option is not enabled (this is the secure default). Run the following command to review the sudoers configuration:

```
# grep -rPi -- '^\h*Defaults\h+([^#]+,\h*)?visiblepw\b' /etc/sudoers
/etc/sudoers.d/
```

Nothing should be returned.

Remediation:

Ensure sudo will not prompt for a password on a terminal where it would be displayed. Edit the appropriate sudoers file with `visudo` and remove any line that enables visiblepw, or explicitly disable it:

```
# visudo
```

Add or correct the line:

```
Defaults !visiblepw
```

Default Value:

visiblepw is not enabled (disabled) by default.

References:

1. GRID: NIX-00000420
2. SUDOERS(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.	●	●	●

5.2.10 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.

Impact:

A shorter authentication timeout causes users to re-enter their password more often within a session.

Audit:

Verify that `timestamp_timeout` is not disabled and is not set greater than 15 minutes.

Run the following command to list any configured value:

```
# grep -roP "timestamp_timeout=\K[0-9]*" /etc/sudoers /etc/sudoers.d/
```

If no `timestamp_timeout` is configured in these files, sudo's compiled default of 5 minutes applies, which is within the required maximum of 15 minutes.

Remediation:

Edit the file listed in the audit section with `visudo -f <PATH TO FILE>` and modify the entry `timestamp_timeout=` to 15 minutes or less as per your site policy. The value is in minutes. This particular entry may appear on its own, or on the same line as `env_reset`. See the following two examples:

Example 1:

```
Defaults    env_reset, timestamp_timeout=15
```

Example 2:

```
Defaults    timestamp_timeout=15
Defaults    env_reset
```

Because sudo applies the last matching `Defaults` definition, place this setting so that no later-parsed file in `/etc/sudoers.d` overrides it with a larger value; see this section's overview for how the sudoers files combine.

Default Value:

timestamp_timeout=5

References:

1. GRID: NIX-00000217
2. NIST SP 800-53 :: IA-11
3. https://www.sudo.ws/man/1.9.0/sudoers.man.html

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.	●	●	●

5.2.11 Ensure sudo timestamp_type is configured (Automated)

Profile Applicability:

- Level 1 - Server

Description:

After a successful authentication, sudo caches the credential for a short period. The `timestamp_type` option controls the scope of that cached credential: `tty` (per terminal, the default), `ppid` (per parent process), `kernel`, or `global` (one record shared by all of the user's processes and sessions).

Rationale:

With `timestamp_type=global`, one successful sudo authentication is honored by every concurrent process and session belonging to that user, widening the window in which an unattended or hijacked session can escalate privileges without re-authenticating. Scoping the timestamp to the terminal (the default, `tty`) limits credential reuse to the session in which the user actually authenticated.

Impact:

Users who operate multiple terminals may be prompted to authenticate in each terminal rather than once for all sessions.

Audit:

Verify that sudo's authentication timestamp is not shared across all of a user's sessions - that is, timestamp_type is not set to global (the default tty scope, or any non-global value, is compliant). Run the following command:

```
# grep -rPi -- 'timestamp_type\h*=\h*global' /etc/sudoers /etc/sudoers.d/
```

Nothing should be returned.

Remediation:

Ensure the sudo authentication timestamp is scoped to the terminal rather than shared globally. Edit the appropriate sudoers file with `visudo` and remove any `Defaults timestamp_type=global`, or set a non-global value:

```
# visudo
Defaults timestamp_type=tty
```

Default Value:

tty

References:

1. GRID: NIX-00000422
2. SUDOERS(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.	●	●	●

5.2.12 Ensure access to the su command is restricted (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `su` command allows a user to run a command or shell as another user. The program has been superseded by `sudo`, which allows for more granular control over privileged access. Normally, the `su` command can be executed by any user. By uncommenting the `pam_wheel.so` statement in `/etc/pam.d/su`, the `su` command will only allow users in 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.

Impact:

Restricting the su command to an empty group effectively prevents its use; administrators must use sudo for privilege escalation. Ensure the appropriate users have sudo access configured before applying.

Audit:

Verify that access to the su command is restricted to an empty group via `pam_wheel.so`.

Run the following command:

```
# grep -Pi
'^\h*auth\h+(?:required|requisite)\h+pam_wheel\.so\h+(?:[^#\n\r]+\h+)?((?!\2)
(use_uid\b|group=\H+\b))\h+(?:[^#\n\r]+\h+)?((?!\1)(use_uid\b|group=\H+\b))(\
h+.*)?$' /etc/pam.d/su
```

Verify the output matches:

```
auth required pam_wheel.so use_uid group=<group_name>
```

Run the following command and verify that the group specified in `<group_name>` contains no users:

```
# grep <group_name> /etc/group
```

Verify the output does not contain any users in the relevant group:

```
<group_name>:x:<GID>:
```

Note: The automated assessment confirms the `pam_wheel.so` configuration line only. Confirming that the specified group contains no users is a manual verification step.

Remediation:

Create an empty group that will be specified for use of the `su` command. The group should be named according to site policy.

Example:

```
# groupadd sugroup
```

Add the following line to the `/etc/pam.d/su` file, specifying the empty group:

```
auth required pam_wheel.so use_uid group=sugroup
```

References:

1. 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) provide the modular framework that programs use to authenticate users on Linux, implemented as shared modules configured under /etc/pam.d. On this platform the PAM configuration is managed through authselect, which applies a selected profile to compose the authentication stack rather than having those files edited by hand. This section covers that configuration across three areas: the software packages that supply PAM and its tooling (pam, authselect, and libpwquality), kept at a current version; the authselect profile that manages PAM, and the enabling of the key modules it composes (pam_faillock, pam_pwquality, pam_pwhistory, and pam_unix); and the arguments that tune those modules, covering account lockout, password quality, password history, and the pam_unix password settings.

Together these establish a coherent, centrally managed authentication policy: failed-logon lockout to slow credential guessing, password quality and history requirements to resist weak or reused passwords, and a strong password-hashing configuration. Because the stack is managed by authselect, configuration changes should be made through authselect rather than by editing the effective PAM files directly, since a later profile application overwrites manual edits. Systems that delegate authentication to a centralized identity provider (for example an IdM or Active Directory domain through SSSD) may manage parts of this configuration outside these files and should be assessed accordingly.

5.3.1 Configure PAM software packages

This section provides recommendations for ensuring the packages that provide Pluggable Authentication Modules (PAM) and password-quality checking are installed and kept up to date.

`pam` provides the PAM authentication framework used across login and privilege-escalation services; `authselect` manages the system's active PAM profile; and `libpwquality` provides the password-quality checking used by `pam_pwquality`. Keeping these packages current ensures the modules and options relied on by the recommendations in the remaining subsections are available, and that security fixes are applied.

5.3.1.1 Ensure latest version of pam is installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The pam package provides the Pluggable Authentication Modules (PAM) framework and the core modules used by the system's authentication, login, and password-change services.

Rationale:

To ensure the system has access to the PAM module functionality used by this Benchmark, the installed version of pam should be 1.5.1-14 or later.

Audit:

Run the following command to verify that pam version 1.5.1-14 or later is installed:

```
# rpm -q pam
```

The version returned should be 1.5.1-14 or later.

Example output:

```
pam-1.5.1-14.el9.x86_64
```

Remediation:

Run the following command to install or update pam to the latest version:

```
# dnf install pam
```

Default Value:

pam is installed by default; the package version present depends on the installed operating system minor release.

References:

1. GRID: NIX-00000222

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			

5.3.1.2 Ensure latest version of authselect is installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Authselect is a utility that simplifies the configuration of user authentication. Authselect offers ready-made profiles that can be universally used with all modern identity management systems.

You can create and deploy a custom profile by customizing one of the default profiles: the sssd, winbind, or the nis profile. This is particularly useful if modifying a ready-made authselect profile is not enough for your needs. When you deploy a custom profile, the profile is applied to every user logging into the given host. This is the recommended method, so that the existing profiles can remain unmodified.

Authselect only modifies files in the following locations:

/etc/nsswitch.conf /etc/pam.d/* /etc/dconf/db/distro.d/*

Rationale:

To ensure the system has access to the authselect functionality used by this Benchmark, the installed version of authselect should be 1.2.6-1 or later.

Impact:

If local site customizations have been made to an authselect default or custom profile created with the `--symlink-pam` option, these customizations may be over-written by updating authselect.

WARNING:

Do not use authselect if:

- your host is part of Linux Identity Management. Joining your host to an IdM domain with the ipa-client-install command automatically configures SSSD authentication on your host.
- Your host is part of Active Directory via SSSD. Calling the realm join command to join your host to an Active Directory domain automatically configures SSSD authentication on your host.

It is **not** recommended to change the authselect profiles configured by ipa-client-install or realm join. If you need to modify them, display the current settings before making any modifications, so you can revert back to them if necessary

Audit:

Run the following command to verify that authselect version 1.2.6-1 or later is installed:

```
# rpm -q authselect
```

The version returned should be 1.2.6-1 or later.

Example output:

```
authselect-1.2.6-1.el9.x86_64
```

Remediation:

Run the following command to install or update authselect to the latest version:

```
# dnf install authselect
```

Default Value:

authselect is installed by default; the package version present depends on the installed operating system minor release.

References:

1. GRID: NIX-00000381

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			

5.3.1.3 Ensure latest version of libpwquality 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 password requirements reduce the risk of successful password guessing and brute-force attacks. Ensuring a current version of libpwquality is installed makes the password-quality checking features used by this Benchmark available.

Audit:

Run the following command to verify that libpwquality version 1.4.4-8 or later is installed:

```
# rpm -q libpwquality
```

The version returned should be 1.4.4-8 or later.

Example output:

```
libpwquality-1.4.4-8.el9.x86_64
```

Remediation:

Run the following command to install or update libpwquality to the latest version:

```
# dnf install libpwquality
```

Default Value:

libpwquality is installed by default; the package version present depends on the installed operating system minor release.

References:

1. GRID: NIX-00000221

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			

5.3.2 Configure authselect

This section verifies that key PAM authentication features - `pam_faillock` (account lockout after failed authentication), `pam_pwquality` (password-complexity checking), `pam_pwhistory` (password-reuse prevention), and `pam_unix` (standard local password authentication) - are in use and are configured through authselect.

authselect is the recommended method for managing PAM on this platform. It renders the effective `/etc/pam.d/system-auth` and `/etc/pam.d/password-auth` files (along with `nsswitch` and related `dconf` settings) from a single managed profile, giving one supported source of truth that is reproducible and upgrade-safe. Manual edits to the generated files are overwritten the next time `authselect apply-changes` runs, so ad hoc changes made directly to `/etc/pam.d` are not durable. `authselect check` verifies that the deployed files are consistent with the active profile, which detects drift or tampering.

An authselect profile is a set of template files. The default profiles ship read-only under `/usr/share/authselect`; to persist customizations, create a custom profile (`authselect create-profile <name> -b <base>`), which is stored under `/etc/authselect/custom/<name>`, select it (`authselect select custom/<name>`), enable the optional features required by this section (`authselect enable-feature with-faillock`, `authselect enable-feature with-pwhistory`; `pam_pwquality` and `pam_unix` are part of the base profile and are not feature toggles), and run `authselect apply-changes` to regenerate the effective `/etc/pam.d` files. Because the default profile templates carry the faillock and pwhistory lines unconditionally inside `{include if "with-faillock"}` / `{include if "with-pwhistory"}` guards, the presence of a module line in a template does not prove the feature is enabled; the audits in this section therefore inspect the effective `/etc/pam.d/system-auth` and `/etc/pam.d/password-auth` files that PAM enforces at runtime.

Restoring a valid authselect configuration

If `authselect check` reports the configuration is not valid, the deployed `/etc/pam.d` files have drifted from the active profile (typically from manual edits). Restore a consistent, authselect-managed state:

1. If an authselect profile is already active, regenerate the files from it:

```
# authselect apply-changes
```

This overwrites any out-of-band edits to the generated files.

2. If no authselect profile is active, select one (with the required features) and apply it:

```
# authselect select <profile> <features>
# authselect apply-changes
```

Use `--force` if existing files must be overwritten.

Systems where PAM is managed outside authselect

A system whose PAM stack is managed by another method (for example, direct edits to `/etc/pam.d` or a third-party configuration manager) will report `authselect check` as not valid and will be assessed as non-compliant by design. Where local site policy requires managing PAM by another method, ensure the covered features (`pam_faillock`, `pam_pwquality`, `pam_pwhistory`, and `pam_unix`) are enabled through that method, then formally document and risk-accept the deviation.

Identity Management and Active Directory-joined hosts

On hosts joined to a Linux Identity Management (IdM) domain or to Active Directory, authselect is configured automatically by the domain-join tooling - `ipa-client-install` for IdM, and `realm join` (via SSSD) for Active Directory. Do not manually alter the authselect profile these tools configure, as doing so can disrupt domain authentication. If a change is required - for example, to enable a feature covered by this section - first record the current settings with `authselect current` so the configuration can be restored, then make the change through authselect and run `authselect apply-changes`.

5.3.2.1 Ensure PAM is managed by an active authselect profile (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

authselect renders and manages the system's PAM configuration from a single active profile. When PAM is managed by authselect, the effective /etc/pam.d/system-auth and /etc/pam.d/password-auth files are generated from that profile and their consistency can be verified, providing one supported, reproducible source of truth for authentication configuration.

Rationale:

A custom profile is required to customize many of the pam options.

Modifications made to a default profile may be overwritten during an update.

When you deploy a profile, the profile is applied to every user logging into the given host

Impact:

Converting a system to authselect management regenerates /etc/pam.d/system-auth and /etc/pam.d/password-auth from the active profile, overwriting any manual edits to those files.

Audit:

Verify PAM is managed by an active authselect profile and the configuration is consistent with it.

Run the following command to verify an authselect profile is active:

```
# authselect current
```

Example output:

```
Profile ID: sssd
Enabled features:
- with-sudo
```

A returned Profile ID indicates an authselect profile is active.

Run the following command to verify the configuration is consistent with the active profile:

```
# authselect check
```

Example output:

```
Current configuration is valid.
```

A valid result indicates PAM is managed by authselect and has not been modified outside it.

Note: The `/etc/pam.d/system-auth` and `/etc/pam.d/password-auth` files are generated by authselect and are the configuration enforced at runtime. If `authselect check` reports the configuration is not valid, PAM has been modified outside authselect.

Remediation:

Configure the system to manage PAM through an authselect profile.

Select an authselect profile (a default profile such as sssd, or a custom profile) and apply it:

```
# authselect select sssd
# authselect apply-changes
```

Use `# authselect select <profile> --force` if existing files must be overwritten (for example, on a system not previously managed by authselect).

If `authselect check` reports the configuration is not valid, restore a valid configuration as described in the section overview.

References:

1. GRID: NIX-00000425
2. authselect(8)
3. https://docs.oracle.com/en/operating-systems/oracle-linux/9/userauth/OL9-USERAUTH.pdf
4. https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html/configuring_authentication_and_authorization_in_rhel/configuring-user-authentication-using-authselect_configuring-authentication-and-authorization-in-rhel#creating-and-deploying-your-own-authselect-profile_configuring-user-authentication-using-authselect

Additional Information:

With the option `--base-on=BASE-ID` (or `-b=BASE-ID`) a new profile is based on the profile named BASE-ID.

The base profile location is determined with these steps:

1. If BASE-ID starts with the prefix `custom/` it is a custom profile.
2. Try if BASE-ID is found in vendor profiles.
3. Try if BASE-ID is found in default profiles.
4. Return an error.

The authselect option `--force` (or `-f`) causes authselect to write changes even when the previous configuration was not created by authselect (for example, files created by another tool or by manual changes). This option automatically backs up system files before writing any change unless `--nobackup` is set.

Example:

```
# authselect select custom/custom-profile --backup=PAM_CONFIG_BACKUP --force
```

Note: The PAM and authselect packages must be versions `pam-1.5.1-14` and `authselect-1.2.6-1` or newer.

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			

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:

Enabling `pam_faillock` introduces account-lockout behavior: after the configured number of consecutive failed authentication attempts an account is locked for the configured interval. An aggressive threshold can lock out legitimate users and become a self-inflicted denial-of-service vector, so confirm the lockout thresholds are appropriate for the environment and be familiar with the unlock procedure (`faillock --user <name> --reset`).

Audit:

Verify the pam_faillock module is enabled in the effective PAM configuration on an authselect-managed system.

1. Run the following command to verify the configuration is managed by authselect and is consistent with the active profile:

```
# authselect check
```

Example output:

```
Current configuration is valid.
```

2. Run the following command to verify the `pam_faillock.so` lines are present in the effective files PAM reads at runtime:

```
# grep -P -- '\bpam_faillock\.so\b' /etc/pam.d/system-auth
/etc/pam.d/password-auth
```

An `auth` preauth line, an `auth` authfail line, and an `account` line for `pam_faillock.so` in both files indicate the module is enabled.

Example output:

```
/etc/pam.d/system-auth:auth        required      pam_faillock.so preauth
/etc/pam.d/system-auth:auth        required      pam_faillock.so authfail
/etc/pam.d/system-auth:account     required      pam_faillock.so
/etc/pam.d/password-auth:auth      required      pam_faillock.so preauth
/etc/pam.d/password-auth:auth      required      pam_faillock.so authfail
/etc/pam.d/password-auth:account   required      pam_faillock.so
```

Note: `/etc/pam.d/system-auth` and `/etc/pam.d/password-auth` are generated by authselect and are the configuration enforced at runtime. If `authselect check` reports the configuration is not valid, PAM has been modified outside authselect and these files are not authoritative.

Remediation:

On a system already managed by an authselect profile, run the following commands to enable the faillock feature and apply the changes:

```
# authselect enable-feature with-faillock
# authselect apply-changes
```

authselect regenerates `/etc/pam.d/system-auth` and `/etc/pam.d/password-auth` with the pam_faillock.so lines.

If the system is not yet managed by authselect, or `authselect check` reports the configuration is not valid, first establish or restore the profile as described in the section overview, then run the commands above.

References:

1. GRID: NIX-00000223
2. pam_faillock(8) - Linux man page
3. NIST SP 800-53 :: AC-7 a
4. NIST SP 800-53A :: AC-7.1 (ii)
5. faillock(8) - Linux man page

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			

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:

`pam_pwquality` enforces password-complexity requirements when users set or change passwords. Overly strict complexity settings can frustrate users, increase help-desk password-reset volume, and encourage insecure workarounds, so complexity thresholds should be balanced against usability.

Audit:

Verify the pam_pwquality module is enabled in the effective PAM configuration on an authselect-managed system.

1. Run the following command to verify the configuration is managed by authselect and is consistent with the active profile:

```
# authselect check
```

Example output:

```
Current configuration is valid.
```

2. Run the following command to verify the `pam_pwquality.so` line is present in the effective files PAM reads at runtime:

```
# grep -P -- '\bpam_pwquality\.so\b' /etc/pam.d/system-auth
/etc/pam.d/password-auth
```

A `password` (`requisite` or `required`) line for `pam_pwquality.so` in both files indicates the module is enabled.

Example output:

```
/etc/pam.d/system-auth:password    requisite     pam_pwquality.so
local_users_only
/etc/pam.d/password-auth:password  requisite     pam_pwquality.so
local_users_only
```

Note: `/etc/pam.d/system-auth` and `/etc/pam.d/password-auth` are generated by authselect and are the configuration enforced at runtime. If `authselect check` reports the configuration is not valid, PAM has been modified outside authselect and these files are not authoritative.

Remediation:

pam_pwquality is part of the default authselect profiles, so a system using a default profile already provides it.

On a system already managed by an authselect profile, ensure the active profile includes pam_pwquality:

- The default profiles include it; no change is required.
- If a custom profile has removed it, add the following line to the custom profile's `system-auth` and `password-auth` templates in `/etc/authselect/custom/<profile>/`:

```
password    requisite    pam_pwquality.so
```

Then apply the configuration:

```
# authselect apply-changes
```

If the system is not yet managed by authselect, or `authselect check` reports the configuration is not valid, first establish or restore the profile as described in the section overview.

Note: pam_pwquality is a base-profile module - there is no `with-pwquality` feature to enable (unlike `with-faillock` and `with-pwhistory`).

References:

1. GRID: NIX-00000226
2. NIST SP 800-53A :: CM-6.1 (iv)
3. NIST SP 800-53 :: CM-6 b

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.		●	●

5.3.2.4 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.

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.

Impact:

`pam_pwhistory` prevents reuse of recently used passwords by recording password history. Users may need more effort to choose a new password that has not been used recently, and a larger remembered-history depth increases that effort.

Audit:

Verify the pam_pwhistory module is enabled in the effective PAM configuration on an authselect-managed system.

1. Run the following command to verify the configuration is managed by authselect and is consistent with the active profile:

```
# authselect check
```

Example output:

```
Current configuration is valid.
```

2. Run the following command to verify the `pam_pwhistory.so` line is present in the effective files PAM reads at runtime:

```
# grep -P -- '\bpam_pwhistory\.so\b' /etc/pam.d/system-auth
/etc/pam.d/password-auth
```

A `password` (`requisite` or `required`) line for `pam_pwhistory.so` in both files indicates the module is enabled.

Example output:

```
/etc/pam.d/system-auth:password    requisite     pam_pwhistory.so use_authtok
/etc/pam.d/password-auth:password  requisite     pam_pwhistory.so use_authtok
```

Note: `/etc/pam.d/system-auth` and `/etc/pam.d/password-auth` are generated by authselect and are the configuration enforced at runtime. If `authselect check` reports the configuration is not valid, PAM has been modified outside authselect and these files are not authoritative.

Remediation:

On a system already managed by an authselect profile, run the following commands to enable the pwhistory feature and apply the changes:

```
# authselect enable-feature with-pwhistory
# authselect apply-changes
```

authselect regenerates `/etc/pam.d/system-auth` and `/etc/pam.d/password-auth` with the pam_pwhistory.so lines.

If the system is not yet managed by authselect, or `authselect check` reports the configuration is not valid, first establish or restore the profile as described in the section overview, then run the commands above.

References:

1. GRID: NIX-00000225
2. 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.		●	●

5.3.2.5 Ensure pam_unix module is enabled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

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.

Rationale:

Requiring users to use authentication make it less likely that an attacker will be able to access the system.

Impact:

`pam_unix` is included in the default authselect profiles and provides standard local password authentication, so ensuring the active profile includes it has no adverse operational impact on a default configuration.

Audit:

Verify the pam_unix module is enabled in the effective PAM configuration on an authselect-managed system.

1. Run the following command to verify the configuration is managed by authselect and is consistent with the active profile:

```
# authselect check
```

Example output:

```
Current configuration is valid.
```

2. Run the following command to verify the `pam_unix.so` lines are present in the effective files PAM reads at runtime:

```
# grep -P -- '\bpam_unix\.so\b' /etc/pam.d/system-auth /etc/pam.d/password-
auth
```

An `auth` line, an `account` line, a `password` line, and a `session` line for `pam_unix.so` in both files indicate the module is enabled.

Example output:

```
/etc/pam.d/system-auth:auth        sufficient    pam_unix.so
/etc/pam.d/system-auth:account     required      pam_unix.so
/etc/pam.d/system-auth:password    sufficient    pam_unix.so
/etc/pam.d/system-auth:session     required      pam_unix.so
/etc/pam.d/password-auth:auth      sufficient    pam_unix.so
/etc/pam.d/password-auth:account   required      pam_unix.so
/etc/pam.d/password-auth:password  sufficient    pam_unix.so
/etc/pam.d/password-auth:session   required      pam_unix.so
```

Note: `/etc/pam.d/system-auth` and `/etc/pam.d/password-auth` are generated by authselect and are the configuration enforced at runtime. If `authselect check` reports the configuration is not valid, PAM has been modified outside authselect and these files are not authoritative.

Remediation:

pam_unix is part of the default authselect profiles, so a system using a default profile already provides it.

On a system already managed by an authselect profile, ensure the active profile includes pam_unix:

- The default profiles include it; no change is required.
- If a custom profile has removed it, restore its `auth`, `account`, `password`, and `session` pam_unix.so entries to the custom profile's `system-auth` and `password-auth` templates in `/etc/authselect/custom/<profile>/` (copy them from the corresponding default profile under `/usr/share/authselect/default/<profile>/`).

Then apply the configuration:

```
# authselect apply-changes
```

If the system is not yet managed by authselect, or `authselect check` reports the configuration is not valid, first establish or restore the profile as described in the section overview.

Note: pam_unix is a base-profile module - there is no `with-unix` feature to enable (unlike `with-faillock` and `with-pwhistory`).

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	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

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.

5.3.3.1 Configure pam_faillock module

This section covers configuring the pam_faillock module's account-lockout behavior - the number of failed authentication attempts before an account is locked, how long a locked account remains locked, and whether the root account is subject to lockout.

faillock.conf provides a way to configure the settings for locking a user after multiple failed authentication attempts, and is read by the pam_faillock module.

The pam_faillock module must first be enabled through the active authselect profile (via the profile's faillock feature) for these settings to take effect. On an authselect-managed system, authselect apply-changes regenerates the PAM configuration from the profile template, including any arguments placed inline on the pam_faillock.so lines - so inline arguments are not durable. Because authselect does not manage /etc/security/faillock.conf, it is the durable and preferred location for the lockout values configured in this section.

faillock.conf uses a simple name = value format, with comments starting with #. Whitespace at the start of a line, at the end of a line, and around the = sign is ignored.

Options relevant to this section:

- deny=<n> - Lock the account after <n> consecutive failed authentication attempts within the fail_interval window. A value of 0 disables the lockout. The default is 3.
- fail_interval=<n> - The interval, in seconds, within which the consecutive failures counted by deny must occur. The default is 900 (15 minutes).
- unlock_time=<n> - Re-enable access <n> seconds after the lockout. A value of 0 means the account is not automatically unlocked and must be reset with the faillock(8) command. The default is 600 (10 minutes).
- even_deny_root - Apply the lockout to the root account as well as regular accounts.
- root_unlock_time=<n> - Implies even_deny_root; re-enable access to the root account <n> seconds after lockout. If not set, it defaults to the unlock_time value.

See faillock.conf(5) for the complete set of options.

Example /etc/security/faillock.conf:

```
deny=5
unlock_time=900
even_deny_root
```

Note: A locked account is re-enabled automatically after unlock_time elapses, or can be unlocked immediately by resetting its failure records with the following command:

```
# faillock --user <username> --reset
```

5.3.3.1.1 Ensure password failed attempts lockout is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `pam_faillock` module locks an account after a number of consecutive failed authentication attempts. The `deny` option, set in `/etc/security/faillock.conf`, defines that threshold: the number of consecutive failed authentication attempts after which the account is locked. A value of `0` disables the lockout.

Rationale:

Locking an account after a defined number of consecutive failed authentication attempts limits the number of password guesses an attacker can make in a short period, mitigating brute force and password-spraying attacks. Setting `deny` to a low, non-zero value ensures an account is locked before an attacker can attempt a large number of credentials.

Audit:

Run the following command to verify that the number of failed authentication 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
```

Example output:

```
deny = 5
```

Run the following command to verify that the `deny` argument has not been set, or is set to 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/system-auth
/etc/pam.d/password-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 script to remove the `deny` argument from the `pam_faillock.so` module in the PAM files:

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

{
   for l_pam_file in system-auth password-auth; do
     l_authselect_file="/etc/authselect/$(head -1
/etc/authselect/authselect.conf | grep 'custom/')/$l_pam_file"
     sed -ri
's/(^\s*auth\s+(requisite|required|sufficient)\s+pam_faillock\.so.*)(\s+deny\
s*=\s*\S+)(.*$)/\1\4/' "$l_authselect_file"
   done
   authselect apply-changes
}
```

Default Value:

deny = 3

References:

1. GRID: NIX-00000235
2. faillock.conf(5)
3. pam_faillock(8)
4. NIST SP 800-53 Rev. 5: IA-5

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.		●	●

5.3.3.1.2 Ensure password unlock time is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `unlock_time` option, set in `/etc/security/faillock.conf`, controls how long a locked account remains locked before access is automatically re-enabled. Access is restored `unlock_time` seconds after the account is locked. A value of `0` means the account is never automatically unlocked and access can only be restored by resetting the faillock entries with the `faillock(8)` command.

Rationale:

A nonzero `unlock_time` automatically restores access after the configured interval, balancing protection against brute force attacks with the operational cost of locked-out accounts and reducing administrative overhead. Setting `unlock_time` to `0` maximizes protection by keeping an account locked until an administrator resets it, at the cost of requiring manual recovery.

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
```

Example output:

```
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:

```
# 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/system-
auth /etc/pam.d/password-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 script to remove the `unlock_time` argument from the `pam_faillock.so` module in the PAM files:

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

{
   for l_pam_file in system-auth password-auth; do
     l_authselect_file="/etc/authselect/$(head -1
/etc/authselect/authselect.conf | grep 'custom/')/$l_pam_file"
     sed -ri
's/(^\s*auth\s+(requisite|required|sufficient)\s+pam_faillock\.so.*)(\s+unloc
k_time\s*=\s*\S+)(.*$)/\1\4/' "$l_authselect_file"
   done
   authselect apply-changes
}
```

Default Value:

unlock_time = 600

References:

1. GRID: NIX-00000244
2. faillock.conf(5)
3. pam_faillock(8)
4. NIST SP 800-53 Rev. 5: IA-5

Additional Information:

Note:

- The directory that `pam_faillock` uses to store failure records is usually cleared on system boot, so access is also re-enabled after a reboot. If that is undesirable, set a different tally directory with the `dir` option.
- It is usually undesirable to permanently lock out users, as they can easily become the target of a denial-of-service attack unless usernames are random and kept secret from potential attackers.
- The maximum configurable value for `unlock_time` is `604800` (7 days).

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.		●	●

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` - 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.

Rationale:

By default the root account is exempt from `pam_faillock` lockout, leaving it able to be subjected to unlimited password-guessing attempts. Because root is a high-value target, including it in the lockout with `even_deny_root` applies the same brute force protection to the most privileged account. A finite `root_unlock_time` can be used to balance this protection against the need to recover the root account.

Impact:

Enabling `even_deny_root` causes the root account to be locked after repeated failed authentication attempts, the same as any other account. On a system with no alternate privileged access path (for example, no `sudo`-capable user and no console or single-user recovery), a locked root account can block emergency administration. Before enabling this on remote or headless systems, ensure an alternate recovery path exists, or set a finite `root_unlock_time` so the root account unlocks automatically after the configured interval. Setting `root_unlock_time=0` keeps the root account locked until an administrator resets the faillock entries, which may cause a denial of service to legitimate administrators.

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
```

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/system-auth /etc/pam.d/password-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 script to remove the `even_deny_root` and `root_unlock_time` arguments from the `pam_faillock.so` module in the PAM files:

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

{
   for l_pam_file in system-auth password-auth; do
     l_authselect_file="/etc/authselect/$(head -1
/etc/authselect/authselect.conf | grep 'custom/')/$l_pam_file"
     sed -ri
's/(^\s*auth\s+(.*)\s+pam_faillock\.so.*)(\s+even_deny_root)(.*$)/\1\4/'
"$l_authselect_file"
     sed -ri
's/(^\s*auth\s+(.*)\s+pam_faillock\.so.*)(\s+root_unlock_time\s*=\s*\S+)(.*$)
/\1\4/' "$l_authselect_file"
   done
   authselect apply-changes
}
```

Default Value:

disabled

References:

1. GRID: NIX-00000234
2. faillock.conf(5)
3. pam_faillock(8)
4. NIST SP 800-53 Rev. 5: IA-5

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.		●	●

5.3.3.2 Configure pam_pwquality module

This section covers configuring the pam_pwquality module's password-quality requirements - minimum length, character complexity, the number of changed characters, limits on repeated and sequential characters, dictionary checking, and whether these requirements apply to the root account.

The `pam_pwquality.so` module checks the strength of passwords. It performs checks such as making sure a password is not a dictionary word, that it is a certain length, and that it contains a mix of characters (for example alphabetic, numeric, and other classes). These checks are configurable by any of:

- use of the module arguments on the `pam_pwquality.so` line
- modifying the `/etc/security/pwquality.conf` configuration file
- creating a `.conf` file in the `/etc/security/pwquality.conf.d/` directory

pam_pwquality is enabled through the active authselect profile. Running `authselect apply-changes` regenerates the PAM configuration from the profile, so any pwquality arguments placed inline on the `pam_pwquality.so` lines are not durable - they are removed the next time the profile is applied. For this reason these settings must be set in `/etc/security/pwquality.conf` or in a `.conf` file in `/etc/security/pwquality.conf.d/` to be durable; a value configured only inline will not pass the audit and must be moved to the durable configuration.

pwquality reads the module arguments first, then `/etc/security/pwquality.conf`, then the `*.conf` files in `/etc/security/pwquality.conf.d/` in ASCII order, with the last value read taking precedence. The recommended durable location for these settings is a `.conf` file in `/etc/security/pwquality.conf.d/`.

Example `/etc/security/pwquality.conf.d/50-pwquality.conf` file:

```
minlen = 14
minclass = 4
difok = 2
maxrepeat = 3
maxsequence = 3
dictcheck = 1
enforce_for_root
```

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:

A new password that differs from the previous password in only a few characters can be trivially derived from a known, compromised, or expired password. Requiring that a minimum number of characters differ from the old password prevents users from making minor, predictable variations of a prior password when it is changed.

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
```

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/system-auth /etc/pam.d/password-
auth
```

Nothing should be returned.

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:

```
difok = 2
```

Example:

```
# sed -ri 's/^\s*difok\s*=/# &/' /etc/security/pwquality.conf
# printf '%s\n' "" "difok = 2" >> /etc/security/pwquality.conf.d/50-
pwdifok.conf
```

Run the following script to remove setting `difok` on the `pam_pwquality.so` module in the PAM files:

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

{
   for l_pam_file in system-auth password-auth; do
     l_authselect_file="/etc/authselect/$(head -1
/etc/authselect/authselect.conf | grep 'custom/')/$l_pam_file"
     sed -ri
's/(^\s*password\s+(requisite|required|sufficient)\s+pam_pwquality\.so.*)(\s+
difok\s*=\s*\S+)(.*$)/\1\4/' "$l_authselect_file"
   done
   authselect apply-changes
}
```

Default Value:

difok = 1

References:

1. GRID: NIX-00000240
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.		●	●

5.3.3.2.2 Ensure password length is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `minlen` option sets the minimum acceptable length for a new password. When character-class credits are enabled (the default), each credited class can count as one additional character toward satisfying this minimum. The value cannot be set lower than 6.

Rationale:

The effort required to guess or crack a password, particularly through offline brute-force attacks, grows rapidly as the password lengthens. Enforcing a minimum password length raises the work factor an attacker must expend and reduces the effectiveness of exhaustive guessing.

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
```

Example output:

```
/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/system-auth
/etc/pam.d/password-auth
```

Nothing should be returned.

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:

```
# sed -ri 's/^\s*minlen\s*=/# &/' /etc/security/pwquality.conf
# printf '%s\n' "" "minlen = 14" >> /etc/security/pwquality.conf.d/50-
pwlength.conf
```

Run the following script to remove setting `minlen` on the `pam_pwquality.so` module in the PAM files:

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

{
   for l_pam_file in system-auth password-auth; do
     l_authselect_file="/etc/authselect/$(head -1
/etc/authselect/authselect.conf | grep 'custom/')/$l_pam_file"
     sed -ri
's/(^\s*password\s+(requisite|required|sufficient)\s+pam_pwquality\.so.*)(\s+
minlen\s*=\s*[0-9]+)(.*$)/\1\4/' "$l_authselect_file"
   done
   authselect apply-changes
}
```

Default Value:

minlen = 8

References:

1. GRID: NIX-00000238
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.		●	●

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:

Requiring a password to draw from multiple classes of characters - digits, uppercase letters, lowercase letters, and other characters - enlarges the pool of possible passwords an attacker must search. A larger keyspace increases the password's resistance to guessing and brute-force attacks.

Impact:

CIS password complexity requirements may differ from those mandated by local site policy or other frameworks; where they conflict, adherence to site-specific policy is imperative.

Audit:

Run the following command to verify that complexity conforms to local site policy:

```
# grep -Psi -- '^\h*(minclass|[dulo]credit)\b' /etc/security/pwquality.conf
/etc/security/pwquality.conf.d/*.conf
```

Example output:

```
/etc/security/pwquality.conf.d/50-pwcomplexity.conf:minclass = 4
  -- AND/OR --
/etc/security/pwquality.conf.d/50-pwcomplexity.conf:dcredit = -1
/etc/security/pwquality.conf.d/50-pwcomplexity.conf:ucredit = -1
/etc/security/pwquality.conf.d/50-pwcomplexity.conf:ocredit = -1
/etc/security/pwquality.conf.d/50-pwcomplexity.conf:lcredit = -1
```

Run the following command to verify that:

- `minclass` is not set to less than `4`
- `dcredit`, `ucredit`, `lcredit`, and `ocredit` are not set to `0` or greater

```
# grep -Psi --
'^\h*password\h+(requisite|required|sufficient)\h+pam_pwquality\.so\h+([^#\n\
r]+\h+)?(minclass=[0-3]|[dulo]credit=[^-]\d*)\b' /etc/pam.d/system-auth
/etc/pam.d/password-auth
```

Nothing should be returned.

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:

- `minclass = 4`

--AND/OR--

- `dcredit = -_N_`
- `ucredit = -_N_`
- `ocredit = -_N_`
- `lcredit = -_N_`

Example:

```
# sed -ri 's/^\s*minclass\s*=/# &/' /etc/security/pwquality.conf
# printf '\n%s' "minclass = 4" >> /etc/security/pwquality.conf.d/50-
pwcomplexity.conf
```

--AND/OR--

```
# sed -ri 's/^\s*[dulo]credit\s*=/# &/' /etc/security/pwquality.conf
# printf '%s\n' "" "dcredit = -1" "ucredit = -1" "ocredit = -1" "lcredit = -
1" > /etc/security/pwquality.conf.d/50-pwcomplexity.conf
```

Run the following script to remove setting `minclass`, `dcredit`, `ucredit`, `lcredit`, and `ocredit` on the `pam_pwquality.so` module in the PAM files

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

{
   for l_pam_file in system-auth password-auth; do
     l_authselect_file="/etc/authselect/$(head -1
/etc/authselect/authselect.conf | grep 'custom/')/$l_pam_file"
     sed -ri
's/(^\s*password\s+(requisite|required|sufficient)\s+pam_pwquality\.so.*)(\s+
minclass\s*=\s*\S+)(.*$)/\1\4/' "$l_authselect_file"
     sed -ri
's/(^\s*password\s+(requisite|required|sufficient)\s+pam_pwquality\.so.*)(\s+
dcredit\s*=\s*\S+)(.*$)/\1\4/' "$l_authselect_file"
     sed -ri
's/(^\s*password\s+(requisite|required|sufficient)\s+pam_pwquality\.so.*)(\s+
ucredit\s*=\s*\S+)(.*$)/\1\4/' "$l_authselect_file"
     sed -ri
's/(^\s*password\s+(requisite|required|sufficient)\s+pam_pwquality\.so.*)(\s+
lcredit\s*=\s*\S+)(.*$)/\1\4/' "$l_authselect_file"
     sed -ri
's/(^\s*password\s+(requisite|required|sufficient)\s+pam_pwquality\.so.*)(\s+
ocredit\s*=\s*\S+)(.*$)/\1\4/' "$l_authselect_file"
   done
   authselect apply-changes
}
```

Note: CIS password complexity requirements may differ from those mandated by local site policy or other frameworks; where they conflict, adherence to site-specific policy is imperative.

Default Value:

minclass = 0

dcredit = 0

ucredit = 0

ocredit = 0

lcredit = 0

References:

1. GRID: NIX-00000232
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.		●	●

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:

Long runs of a single repeated character are low in entropy and are among the first patterns tried by password-guessing tools. Limiting the number of consecutive identical characters allowed in a password forces greater variety and increases the difficulty of guessing it.

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
```

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/system-auth
/etc/pam.d/password-auth
```

Nothing should be returned.

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:

```
# sed -ri 's/^\s*maxrepeat\s*=/# &/' /etc/security/pwquality.conf
# printf '%s\n' "" "maxrepeat = 3" >> /etc/security/pwquality.conf.d/50-
pwrepeat.conf
```

Run the following script to remove setting `maxrepeat` on the `pam_pwquality.so` module in the PAM files:

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

{
   for l_pam_file in system-auth password-auth; do
     l_authselect_file="/etc/authselect/$(head -1
/etc/authselect/authselect.conf | grep 'custom/')/$l_pam_file"
     sed -ri
's/(^\s*password\s+(requisite|required|sufficient)\s+pam_pwquality\.so.*)(\s+
maxrepeat\s*=\s*\S+)(.*$)/\1\4/' "$l_authselect_file"
   done
   authselect apply-changes
}
```

Default Value:

maxrepeat = 0

References:

1. GRID: NIX-00000243
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.		●	●

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:

Monotonic character sequences such as 12345, abcde, or a keyboard walk are predictable and are commonly attempted by password-guessing tools. Limiting the maximum length of such a sequence in a new password removes an easily guessed construction pattern.

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
```

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:

```
# 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/system-auth
/etc/pam.d/password-auth
```

Nothing should be returned.

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:

```
# sed -ri 's/^\s*maxsequence\s*=/# &/' /etc/security/pwquality.conf
# printf '%s\n' "" "maxsequence = 3" >> /etc/security/pwquality.conf.d/50-
pwmaxsequence.conf
```

Run the following script to remove setting `maxsequence` on the `pam_pwquality.so` module in the PAM files:

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

{
   for l_pam_file in system-auth password-auth; do
     l_authselect_file="/etc/authselect/$(head -1
/etc/authselect/authselect.conf | grep 'custom/')/$l_pam_file"
     sed -ri
's/(^\s*password\s+(requisite|required|sufficient)\s+pam_pwquality\.so.*)(\s+
maxsequence\s*=\s*\S+)(.*$)/\1\4/' "$l_authselect_file"
   done
   authselect apply-changes
}
```

Default Value:

maxsequence = 0

References:

1. GRID: NIX-00000239
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.		●	●

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:

Passwords based on dictionary words are among the first candidates attempted by password-cracking tools. Checking a new password against the cracklib dictionary and rejecting dictionary-based passwords defeats dictionary attacks and reduces the likelihood of a password being successfully guessed.

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
```

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/system-auth /etc/pam.d/password-auth
```

Nothing should be returned.

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
```

Run the following script to remove setting `dictcheck` on the `pam_pwquality.so` module in the PAM files:

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

{
   for l_pam_file in system-auth password-auth; do
     l_authselect_file="/etc/authselect/$(head -1
/etc/authselect/authselect.conf | grep 'custom/')/$l_pam_file"
     sed -ri
's/(^\s*password\s+(requisite|required|sufficient)\s+pam_pwquality\.so.*)(\s+
dictcheck\s*=\s*\S+)(.*$)/\1\4/' "$l_authselect_file"
   done
   authselect apply-changes
}
```

Default Value:

dictcheck = 1

References:

1. GRID: NIX-00000233
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.		●	●

5.3.3.2.7 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.

Rationale:

By default the pam_pwquality checks are not enforced when the root user sets a password: a failed quality check produces only a warning and the password is still accepted. Because the root account is the highest-value target on the system, allowing its password to be set without meeting the quality requirements applied to every other account leaves a significant gap. Enforcing password quality for root closes that gap.

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
```

Example output:

```
/etc/security/pwquality.conf.d/50-pwroot.conf:enforce_for_root
```

Remediation:

Edit or add the following line in a `*.conf` file in `/etc/security/pwquality.conf.d` or in `/etc/security/pwquality.conf`:

Example:

```
# printf '%s\n' "" "enforce_for_root" >> /etc/security/pwquality.conf.d/50-
pwroot.conf
```

Default Value:

disabled

References:

1. GRID: NIX-00000242
2. pwquality.conf(5)
3. pam_pwquality(8)
4. NIST SP 800-53 Rev. 5: IA-5

Additional Information:

Note: The root user is not asked for an old password, so the checks that compare the old and new password are not performed for the root user.

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.		●	●

5.3.3.3 Configure pam_pwhistory module

This section covers configuring the pam_pwhistory module's password-history requirements - how many previous passwords are remembered, whether the history check applies to the root account, and that the module reuses the new password token from the PAM stack.

pam_pwhistory is a PAM module that remembers the last passwords set for each user in order to enforce password-change history and keep users from alternating between the same passwords too frequently. The previous password hashes are saved in the /etc/security/opasswd file.

Note: pam_pwhistory 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. These settings apply only to accounts configured on the local system.

The options relevant to this section are:

- remember=<N> - the number of previous passwords saved for each user and checked against on a password change.
- enforce_for_root - if set, the history check is enforced for the root user as well.
- use_authtok - when changing a password, use the new password already provided by a previously stacked password module rather than prompting again.

The options are described in the pwhistory.conf(5) and pam_pwhistory(8) manual pages.

Enablement and durability (authselect)

pam_pwhistory is enabled through the active authselect profile (see the *Configure authselect* section). Running authselect apply-changes regenerates the PAM configuration files (/etc/pam.d/system-auth and /etc/pam.d/password-auth) from the active profile, so arguments placed inline on the pam_pwhistory.so lines are not durable unless they are part of the profile.

- remember and enforce_for_root are pwhistory.conf options. Set them in /etc/security/pwhistory.conf, which is the durable location. A value set only inline on the pam_pwhistory.so module line will be removed the next time authselect apply-changes runs and will not pass the audit.
- use_authtok is a module-line argument, not a pwhistory.conf option. It must be present on the pam_pwhistory.so lines, so its durable home is the active authselect profile's template (a custom profile). Add it to the profile's system-auth and password-auth templates and run authselect apply-changes.

Note: Options specified on the module command line override the values from the `pwhistory.conf` configuration file. When the same option is read more than once, the last value read takes precedence.

The `/etc/security/pwhistory.conf` file uses a simple `name = value` format, with comments starting with the `#` character. It is a single file; there is no drop-in configuration directory for `pwhistory.conf`.

An example `/etc/security/pwhistory.conf`:

```
remember = 24
enforce_for_root
```

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 users' old passwords and can be checked to ensure that users are not recycling recent passwords. The number of passwords remembered is set by the `remember` argument of the `pam_pwhistory` module.

- remember=<N> - `<N>` is the number of old passwords to remember

Rationale:

Remembering a number of previously used passwords and checking new passwords against them prevents users from cycling back to a recently used - and possibly compromised or expired - password. This defeats password reuse and makes it less likely that an attacker can reuse or guess a password the account has used before.

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
```

Example output:

```
remember = 24
```

Run the following command to verify that the remember option is not set to less than `24` on the `pam_pwhistory.so` module in `/etc/pam.d/password-auth` and `/etc/pam.d/system-auth`:

```
# grep -Pi --
'^\h*password\h+(requisite|required|sufficient)\h+pam_pwhistory\.so\h+([^#\n\
r]+\h+)?remember=(2[0-3]|1[0-9]|[0-9])\b' /etc/pam.d/system-auth
/etc/pam.d/password-auth
```

Nothing should be returned.

Note: The `remember` value must be set in `/etc/security/pwhistory.conf` to be durable. A value set only on the `pam_pwhistory.so` module line is not a passing configuration: `authselect apply-changes` regenerates the PAM files and removes inline arguments, so an inline-only value would not persist.

Remediation:

Edit or add the following line in `/etc/security/pwhistory.conf`:

```
remember = 24
```

Run the following script to remove the `remember` argument from the `pam_pwhistory.so` module in `/etc/pam.d/system-auth` and `/etc/pam.d/password-auth`:

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

{
   for l_pam_file in system-auth password-auth; do
     l_authselect_file="/etc/authselect/$(head -1
/etc/authselect/authselect.conf | grep 'custom/')/$l_pam_file"
     sed -ri
's/(^\s*password\s+(requisite|required|sufficient)\s+pam_pwhistory\.so.*)(\s+
remember\s*=\s*\S+)(.*$)/\1\4/' "$l_authselect_file"
   done
   authselect apply-changes
}
```

Default Value:

Not set

References:

1. GRID: NIX-00000237
2. NIST SP 800-53 Rev. 5: IA-5(1)
3. pwhistory.conf(5)
4. pam_pwhistory(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.		●	●

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:

By default the password-history check does not apply to the root account, so root can be changed back to a recently used password even when history enforcement is configured for other users. Enabling `enforce_for_root` subjects root's password changes to the same history check. Because the root account is a high-value target, preventing reuse of a recent - and possibly compromised - root password closes that gap.

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
```

Example output:

```
enforce_for_root
```

Remediation:

Edit or add the following line in `/etc/security/pwhistory.conf`:

```
enforce_for_root
```

Default Value:

disabled

References:

1. GRID: NIX-00000236
2. NIST SP 800-53 Rev. 5: IA-5
3. pwhistory.conf(5)
4. pam_pwhistory(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.		●	●

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` directs `pam_pwhistory` to use the new password token already supplied by an earlier module in the password stack instead of prompting for it again. This ensures the password actually being set is the one checked against the password history. Without it, the history check may be applied to a different value or skipped, weakening enforcement and allowing a reused password to slip through.

Audit:

Run the following command to verify that `use_authtok` is set on the `pam_pwhistory.so` module lines in the password stack:

```
# grep -P --
'^\h*password\h+([^#\n\r]+)\h+pam_pwhistory\.so\h+([^#\n\r]+\h+)?use_authtok\
b' /etc/pam.d/{password,system}-auth
```

Output should be similar to:

```
/etc/pam.d/password-auth:password    required    pam_pwhistory.so use_authtok

/etc/pam.d/system-auth:password    required    pam_pwhistory.so use_authtok
```

Verify that the lines include `use_authtok`

Remediation:

On an authselect-managed system, `use_authtok` is a module-line argument and must be present on the `pam_pwhistory.so` line in the active profile's password stack. Set it in the active custom profile's templates so that it persists across `authselect apply-changes`.

Edit the active custom profile's `system-auth` and `password-auth` templates under `/etc/authselect/custom/<profile>/` so that the `pam_pwhistory.so` line in the password stack includes `use_authtok`:

```
password    requisite    pam_pwhistory.so    use_authtok
```

Apply the changes to regenerate the PAM configuration:

```
# authselect apply-changes
```

Note: If the system is not yet using a custom authselect profile, first create one from the active base profile and select it, then edit its templates as above and apply:

```
# authselect create-profile <profile> -b <current-base>
# authselect select custom/<profile> <features>
```

See the section overview for the authselect profile mechanics.

References:

1. GRID: NIX-00000224
2. NIST SP 800-53 Rev. 5: IA-5
3. pwhistory.conf(5)
4. pam_pwhistory(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.		●	●

5.3.3.4 Configure pam_unix module

This section covers the arguments on the pam_unix module: preventing blank-password authentication (nullok), keeping password-history handling off pam_unix (remember), the password hashing algorithm, and reuse of the validated password token (use_authtok).

The pam_unix.so module is the standard Unix authentication module. It uses the standard system libraries to retrieve and set account information and to authenticate users, typically against /etc/passwd and, where shadow passwords are enabled, /etc/shadow.

Unlike the other PAM modules configured in this section, pam_unix has no configuration file. Every one of these settings is an argument on the pam_unix.so line in the PAM stack. pam_unix is enabled through the active authselect profile, and running authselect apply-changes regenerates `/etc/pam.d/system-auth` and `/etc/pam.d/password-auth` from that profile. As a result, arguments edited directly in the `/etc/pam.d` files are not durable; they are removed the next time the profile is applied. The durable source is the active custom authselect profile template, so a change made only by editing `/etc/pam.d/*` directly will be reverted and will not pass the audit. Editing a pam_unix module-line argument requires a custom profile, because the stock profiles under /usr/share/authselect/ are package-owned and must not be edited.

1. Determine the active profile:

```
# authselect current
```

2. If the active profile is a stock profile (for example sssd), create a custom profile from it and select it, including any features the system already had enabled (authselect current lists them):

```
# authselect create-profile custom-profile -b sssd
# authselect select custom/custom-profile
```

3. Set the pam_unix arguments on the pam_unix.so lines in the custom profile's system-auth and password-auth templates under /etc/authselect/custom/<profile>/, then apply the change so the /etc/pam.d files are regenerated:

```
# authselect apply-changes
```

Feature toggles such as without-nullok are enabled with authselect enable-feature and do not require a custom profile.

The password hashing algorithm is also a pam_unix.so argument. yescrypt is the default on this platform and sha512 is also acceptable; both provide a strong one-way hash for stored passwords.

Note: These settings apply only to accounts configured on the local system.

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:

The nullok argument permits authentication for a local account that has a blank or empty password, allowing access to a pam_unix backed service with no password. Removing nullok ensures every account must present a real password to authenticate.

Audit:

Run the following command to verify the nullok argument is not present on any pam_unix.so line in the effective PAM configuration:

```
# grep -Pi --
'^\h*(auth|account|password|session)\h+(requisite|required|sufficient)\h+pam_
unix\.so\h+([^#\n\r]+\h+)?nullok\b' /etc/pam.d/system-auth
/etc/pam.d/password-auth
```

Nothing should be returned.

Remediation:

without-nullok is provided as an authselect feature, so it is enabled on the active profile directly and does not require a custom profile (see the section overview for the authselect profile model).

Enable the feature on the active profile:

```
# authselect enable-feature without-nullok
```

Apply the change so that /etc/pam.d/system-auth and /etc/pam.d/password-auth are regenerated without the nullok argument:

```
# authselect apply-changes
```

Default Value:

Not set. By default the pam_unix.so module does not permit authentication for an account with a blank password (nullok is absent).

References:

1. GRID: NIX-00000227
2. NIST SP 800-53A :: CM-6.1 (iv)
3. NIST SP 800-53 :: CM-6 b
4. pam_unix(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.		●	●

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 `sha512` hash algorithm to be used.

Rationale:

Password history is enforced through the pam_pwhistory module, which records prior passwords using the strong hash configured on pam_unix. The remember argument on pam_unix is the legacy mechanism: it stores prior passwords in /etc/security/opasswd using the weak MD5 hash. It should not be set on pam_unix so that password history relies on the stronger pam_pwhistory implementation.

Audit:

Run the following command to verify the remember argument is not present on any pam_unix.so line in the effective PAM configuration:

```
# grep -Pi --
'^\h*password\h+([^#\n\r]+\h+)?pam_unix\.so\h+([^#\n\r]+\h+)?remember='
/etc/pam.d/system-auth /etc/pam.d/password-auth
```

Nothing should be returned.

Remediation:

The pam_unix arguments are set on the pam_unix.so module lines in the active authselect profile; editing /etc/pam.d/system-auth or /etc/pam.d/password-auth directly is not durable, because authselect regenerates those files (see the section overview for the authselect profile model).

Remove the remember argument so password history is handled by pam_pwhistory rather than pam_unix:

1. Ensure the system is using a custom authselect profile. If it is still on a stock profile, create and select one first, as described in the section overview.
2. On the pam_unix.so lines in the active custom profile templates (/etc/authselect/custom/<profile>/system-auth and /etc/authselect/custom/<profile>/password-auth), remove any remember=N argument. A compliant line reads:

```
password    sufficient    pam_unix.so yescrypt shadow use_authtok
```

3. Apply the change so that /etc/pam.d/system-auth and /etc/pam.d/password-auth are regenerated without the remember argument:

```
# authselect apply-changes
```

Default Value:

Not set. The pam_unix.so module does not enable password history (remember is absent) by default.

References:

1. GRID: NIX-00000228
2. pam_unix(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.		●	●

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.

Rationale:

Weak or legacy password hashing algorithms such as md5 and DES are fast to compute, so an attacker who obtains the stored password hashes can crack them offline at high speed. A strong algorithm (yescrypt or SHA-512) is deliberately expensive to compute, greatly increasing the effort required to recover local user passwords from the stored hashes.

Audit:

Run the following command to verify a strong password hashing algorithm (sha512 or yescrypt) is set on the pam_unix.so password lines in the effective PAM configuration:

```
# grep -Pi --
'^\h*password\h+(requisite|required|sufficient)\h+pam_unix\.so(\h+[^#\n\r]+)?
\h+(sha512|yescrypt)\b' /etc/pam.d/system-auth /etc/pam.d/password-auth
```

A line should be returned for both files, each including either sha512 or yescrypt.

Example output:

```
/etc/pam.d/password-auth:password sufficient pam_unix.so yescrypt shadow
use_authtok
/etc/pam.d/system-auth:password sufficient pam_unix.so yescrypt shadow
use_authtok
```

Remediation:

The pam_unix arguments are configured through the active authselect profile so that the change is durable; editing /etc/pam.d/system-auth or /etc/pam.d/password-auth directly is not durable, because authselect regenerates those files (see the section overview for the authselect profile model).

Set yescrypt, the default and recommended algorithm on this platform:

1. Ensure the system is using a custom authselect profile. If it is still on a stock profile, create and select one first, as described in the section overview.
2. On the pam_unix.so line in the password stack of the active custom profile templates (`/etc/authselect/custom/<profile>/system-auth` and `/etc/authselect/custom/<profile>/password-auth`), ensure the hashing algorithm is set to yescrypt and remove any other algorithm such as md5, sha256, bigcrypt, or blowfish. A compliant line reads:

```
password    sufficient    pam_unix.so yescrypt shadow use_authtok
```

3. Apply the change so that /etc/pam.d/system-auth and /etc/pam.d/password-auth are regenerated with the strong hashing algorithm:

```
# authselect apply-changes
```

Note: The hashing algorithm applies to passwords set after the change; existing password hashes remain in place until each account's password is next changed.

Systems with a FIPS 140 requirement:

yescrypt is not a FIPS-approved algorithm and must not be used where FIPS compliance is mandated. Use sha512 instead, which is also a strong algorithm, by specifying sha512 in place of yescrypt in step 2 and then applying the change. Such environments additionally enable system-wide FIPS mode, which aligns the system crypto policy with the FIPS-approved algorithm:

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

Default Value:

yescrypt. On this platform the default authselect profile sets yescrypt as the hashing algorithm on the pam_unix.so password lines.

References:

1. GRID: NIX-00000229
2. NIST SP 800-53 Rev. 5: IA-5
3. pam_unix(8)

Additional Information:

Additional module options may be set, recommendation only covers those listed here.

The following command may be used to expire all non-system user ID's immediately and force them to change their passwords on next login. Any system accounts that need to be expired should be carefully done separately by the system administrator to prevent any potential problems.

```
# awk -F: '( $3<'"$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)"' && $1 !=
"nfsnobody" ) { print $1 }' /etc/passwd | xargs -n 1 chage -d 0
```

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.11 Encrypt Sensitive Data at Rest Encrypt sensitive data at rest on servers, applications, and databases containing sensitive data. Storage-layer encryption, also known as server-side encryption, meets the minimum requirement of this Safeguard. Additional encryption methods may include application-layer encryption, also known as client-side encryption, where access to the data storage device(s) does not permit access to the plain-text data.		●	●
v7	16.4 Encrypt or Hash all Authentication Credentials Encrypt or hash with a salt all authentication credentials when stored.		●	●

5.3.3.4.4 Ensure pam_unix includes use_authtok (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The use_authtok argument tells pam_unix, during a password change, to use the new password token already supplied by an earlier module in the password stack rather than prompting for a new one.

Rationale:

With use_authtok, pam_unix commits the password that was already validated by earlier modules in the password stack, such as the quality and history checks, so the password that passed those checks is the one that is stored. Without use_authtok, pam_unix may prompt for and store a different password, bypassing the preceding validation.

Audit:

Run the following command to verify that `use_authtok` is set on the `pam_unix.so` module lines in the password stack:

```
# grep -P --
'^\h*password\h+([^#\n\r]+)\h+pam_unix\.so\h+([^#\n\r]+\h+)?use_authtok\b'
/etc/pam.d/{password,system}-auth
```

Output should be similar to:

```
/etc/pam.d/password-auth:password   sufficient   pam_unix.so sha512 shadow
use_authtok

/etc/pam.d/system-auth:password   sufficient   pam_unix.so sha512 shadow
use_authtok
```

Verify that the lines include `use_authtok`

Remediation:

The pam_unix arguments are set on the pam_unix.so module lines in the active authselect profile; editing /etc/pam.d/system-auth or /etc/pam.d/password-auth directly is not durable, because authselect regenerates those files (see the section overview for the authselect profile model).

Add the use_authtok argument so pam_unix uses the password already validated by the preceding modules:

1. Ensure the system is using a custom authselect profile. If it is still on a stock profile, create and select one first, as described in the section overview.
2. On the pam_unix.so line in the password stack of the active custom profile templates (/etc/authselect/custom/<profile>/system-auth and /etc/authselect/custom/<profile>/password-auth), add use_authtok. A compliant line reads:

```
password    sufficient    pam_unix.so yescrypt shadow use_authtok
```

3. Apply the change so that /etc/pam.d/system-auth and /etc/pam.d/password-auth are regenerated with the use_authtok argument:

```
# authselect apply-changes
```

Default Value:

Not set. The use_authtok argument is not enabled on the pam_unix.so lines by default.

References:

1. GRID: NIX-00000230
2. NIST SP 800-53 Rev. 5: IA-5
3. pam_unix(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.		●	●

5.4 User Accounts and Environment

This section covers the configuration of user accounts and their runtime environment, across three areas: the shadow password suite parameters that govern password aging and account inactivity; the integrity of the root and system accounts (their unique privileged identifiers, controlled access, and hardened environment); and the default environment applied to interactive users, including login shells, idle timeout, and file-creation mode.

Together these reduce the risk of stale or long-lived credentials, of privileged and service accounts being silently duplicated or misused, and of interactive sessions starting from an insecure default. These are host-level account controls and apply across all deployment types, including cloud/CSP images.

5.4.1 Configure shadow password suite parameters

This subsection covers the shadow password suite parameters that govern the local password lifecycle: password expiration (maximum age), the minimum number of days between password changes, the expiration warning period, the password hashing algorithm, inactive-account lockout, and detection of accounts whose last password-change date is in the future.

While most password controls have moved to PAM, these remain in the shadow password suite and exist at two layers: the system-wide defaults (/etc/login.defs, plus /etc/default/useradd for the inactivity lockout), which seed accounts created afterward, and the per-user values in /etc/shadow, which govern existing accounts. Both layers are checked when assessing these controls; existing accounts are updated with the chage command (or usermod where applicable).

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.

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 a value greater than 0, no greater than 365, and conforms to local site policy:

```
# grep -Pi -- '^\h*PASS_MAX_DAYS\h+\d+\b' /etc/login.defs
```

Example output:

```
PASS_MAX_DAYS 365
```

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` and no greater than `365` 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`

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	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

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 affected user should be temporarily changed to `0` via `chage --mindays 0 <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:

```
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	0.0 Explicitly Not Mapped Explicitly Not Mapped			
v7	0.0 Explicitly Not Mapped Explicitly Not Mapped			

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.

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
```

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

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
2. CIS Password Policy Guide
3. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

CIS Controls:

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

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 (used only if `ENCRYPT_METHOD` is unset)

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.

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 SHA512
```

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 SHA512

References:

1. GRID: NIX-00000248
2. NIST SP 800-53 Rev. 5: IA-5

Additional Information:

Note: Under FIPS 140 mode (enabled with `fips-mode-setup --enable`), `ENCRYPT_METHOD` must be `SHA512`; yescrypt is not FIPS-approved. Keep consistent with the pam_unix hashing configuration.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.11 Encrypt Sensitive Data at Rest Encrypt sensitive data at rest on servers, applications, and databases containing sensitive data. Storage-layer encryption, also known as server-side encryption, meets the minimum requirement of this Safeguard. Additional encryption methods may include application-layer encryption, also known as client-side encryption, where access to the data storage device(s) does not permit access to the plain-text data.		●	●
v7	16.4 Encrypt or Hash all Authentication Credentials Encrypt or hash with a salt all authentication credentials when stored.		●	●

5.4.1.5 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.

Audit:

Run the following command and verify `INACTIVE` conforms to site policy (no more than 45 days):

```
# useradd -D | grep INACTIVE

INACTIVE=45
```

Verify all users with a password have Password inactive no more than 45 days after password expires. Run the following command and review the 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 == "" || $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:

```
# 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	5.3 Disable Dormant Accounts Delete or disable any dormant accounts after a period of 45 days of inactivity, where supported.	●	●	●
v7	16.9 Disable Dormant Accounts Automatically disable dormant accounts after a set period of inactivity.	●	●	●

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. The shadow password suite calculates password expiry relative to the last password change date, so an account with a far-future last-change date can permanently bypass password expiration policies: the password never expires regardless of `PASS_MAX_DAYS`.

Impact:

Correcting a user's last password change date to a valid past date may immediately subject the account to the configured `PASS_MAX_DAYS` expiration. If the corrected date already exceeds the maximum password age, the user may be forced to change the password at next login.

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:

```
# chage -d "$(date +%Y-%m-%d)" <user>
```

Force password change on next login (recommended companion):

```
# chage -d 0 <user>
```

Default Value:

Not applicable

References:

1. GRID: NIX-00000250
2. CIS Password Policy Guide
3. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

CIS Controls:

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

5.4.2 Configure root and system accounts and environment

This subsection covers the accounts and environment that carry unrestricted privilege on the system: the uniqueness of the root identity (the sole account with UID 0, the sole account with primary GID 0, and root as the sole group with GID 0), how direct access to the root account is controlled, the integrity of root's PATH and default file-creation mask (umask), and the disposition of the non-login system accounts that own system files and run services.

Two sets of default accounts are expected and are treated as exemptions rather than findings. When verifying that no unexpected account holds primary GID 0, the default system accounts sync, shutdown, halt, and operator are permitted. When verifying that system accounts have no valid login shell, the root, halt, sync, shutdown, and nfsnobody accounts are excluded. The individual recommendations reference this shared exemption set rather than each redefining it.

Some of these controls distinguish a value that is configured in a file from the value that is actually in effect at login. A umask or PATH set in one initialization file can be overridden by a later file or by the login environment, so it is root's effective file-creation mask and effective PATH that determine compliance, not merely the presence of a value in a single startup file.

These are host-level account controls and apply to every deployment type, including physical hosts, virtual machines, and cloud service provider (CSP) images. A hardened base image should already satisfy them, but they must still be assessed on running instances, because provisioning and configuration management can introduce additional privileged or non-login accounts after the image is built.

5.4.2.1 Ensure root is the only UID 0 account (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

On a Linux system, the user ID (UID) `0` is the superuser identity: the kernel grants any account assigned UID `0` complete, unrestricted control of the host, regardless of the account's name. The `root` account holds UID `0` by default.

Rationale:

Limiting UID `0` to the single default `root` account ensures privileged access is accountable and cannot be silently obtained through a second, easily overlooked superuser identity. Administrative access should be performed through an unprivileged account using an approved escalation mechanism, as described in the recommendation covering restriction of the `su` command.

Audit:

Run the following command to verify that root is the only account with UID 0:

```
# awk -F: '($3 == 0) { print $1 }' /etc/passwd
```

Only `root` should be returned.

Remediation:

1. Run the following command to reassign any account other than root that has UID 0 to a new, unused UID (and update ownership of that account's files):

```
# usermod -u <NEW_UID> <username>
```

Example:

```
# usermod -u 1234 bsmith
```

2. If the root account's UID is not 0, run the following command to reset it:

```
# usermod -u 0 root
```

References:

1. GRID: NIX-00000255
2. NIST SP 800-53A :: CM-6.1 (iv)
3. NIST SP 800-53 :: CM-6 b

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.	●	●	●

5.4.2.2 Ensure root is the only GID 0 account (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

An account's primary group is identified by the group ID (GID) in its `/etc/passwd` entry. GID `0` is the `root` group, which the `root` account uses as its primary group by default.

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 that only root and the default system accounts have primary GID 0:

```
# awk -F: '($4 == 0) { print $1 }' /etc/passwd
```

`root` should be returned; the default system accounts `sync`, `shutdown`, `halt`, and `operator` may also be returned.

Remediation:

Run the following command to reassign any account other than root that has primary GID 0 to an appropriate group:

```
# usermod -g <NEW_GID> <username>
```

Example:

```
# usermod -g 100 bsmith
```

If the root account's primary GID is not 0, run the following command to reset it:

```
# usermod -g 0 root
```

References:

1. GRID: NIX-00000254
2. NIST SP 800-53 Rev. 5: CM-1, CM-2, CM-6, CM-7, IA-5

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	5.4 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.	●	●	●

5.4.2.3 Ensure group root is the only GID 0 group (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Groups are defined in `/etc/group`, each identified by a group ID (GID). GID `0` is reserved for the `root` group, the privileged group associated with the superuser.

Rationale:

If a group other than `root` is assigned GID `0`, it silently shares the privileges of the `root` group. Such a duplicate is easy to overlook and undermines accountability for access granted through the root group.

Audit:

Run the following command to verify that root is the only group with GID 0:

```
# awk -F: '($3 == 0) { print $1 }' /etc/group
```

Only `root` should be returned.

Remediation:

Run the following command to reassign any group other than root that has GID 0 to an unused GID:

```
# groupmod -g <NEW_GID> <groupname>
```

Example:

```
# groupmod -g 200 legacyadmins
```

If the root group's GID is not 0, run the following command to reset it:

```
# groupmod -g 0 root
```

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	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.	●	●	●

5.4.2.4 Ensure root account access is controlled (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

There are several ways to gain access to the `root` account directly: authenticating at the console or a login prompt, switching to it with the `su` command, or opening a `root` shell through `sudo`. Direct access to the root account is governed by its authentication state: whether a password is set and whether the account is locked.

Rationale:

If the `root` account has no password set and is not locked, any user who can reach a `root` login prompt or otherwise select the `root` account can gain access without authenticating, and thus gain control over the entire system. Ensuring that the `root` password is set, or that the `root` account is locked so that direct authentication is disabled in favor of an auditable escalation path such as `sudo`, prevents an unauthenticated user from taking control of the system.

Impact:

If there are any automated processes that rely on access to the root account without authentication, they will fail after remediation.

Note: Locking the root account can prevent emergency or rescue (single-user) login. Before locking root, ensure an alternative means to obtain privilege (for example, a sudo-capable administrative account) is in place.

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: PS
- OR -
User: "root" Password is status: LK
```

Note:

- `PS` - Password is set
- `LK` - Password is locked

Remediation:

Run the following command to set a password for the `root` user:

```
# passwd root
```

- OR -

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	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.	●	●	●

5.4.2.5 Ensure root path integrity (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The `PATH` variable for the `root` user determines the directories the shell searches for executables when `root` runs a command that is not given as an absolute path. Its integrity depends on the ownership, permissions, and composition of the directories it lists.

Rationale:

Including the current working directory (.) or other writable directory in `root`'s executable path makes it likely that an attacker can gain superuser access by forcing an administrator operating as `root` to execute a Trojan horse program.

Audit:

Run the following script to review root's PATH; it reports any empty (::), trailing (:), or current-working-directory (.) entry, and any PATH directory that does not exist, is not owned by root, or is writable by group or other:

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

{
   l_pmask="0022"
   l_root_path="$(su - root -c 'printf "%s" "$PATH"')"
   grep -q  -- "::" <<< "$l_root_path" && echo " - root's PATH contains an
empty directory (::)"
   grep -Pq -- ":\h*$" <<< "$l_root_path" && echo " - root's PATH ends with a
trailing colon (:)"
   grep -Pq -- '(^\h*|:)\.(:|\h*$)' <<< "$l_root_path" && echo " - root's
PATH contains the current working directory (.)"
   IFS=":" read -ra a_path_loc <<< "$l_root_path"
   for l_path in "${a_path_loc[@]}"; do
      if [ -d "$l_path" ]; then
         while IFS=: read -r l_fmode l_fown; do
            [ "$l_fown" != "root" ] && echo " - Directory \"$l_path\" is
owned by \"$l_fown\" (should be owned by root)"
            [ $(( l_fmode & l_pmask )) -gt 0 ] && echo " - Directory
\"$l_path\" is mode \"$l_fmode\" (should be 0755 or more restrictive)"
         done <<< "$(stat -Lc '%#a:%U' "$l_path")"
      else
         echo " - \"$l_path\" is not an existing directory"
      fi
   done
}
```

Nothing should be returned.

Remediation:

For each non-compliant item identified by the audit, correct root's PATH as follows.

Edit the shell initialization file that sets root's PATH (for example /root/.bash_profile, /root/.bashrc, or the relevant file under /etc/profile.d/) and remove any current-working-directory (.) entry, any empty entry (::), and any trailing colon (:).

For each directory in root's PATH that is not owned by root, run the following command to set its owner to root:

```
# chown root <directory>
```

For each directory in root's PATH that is writable by group or other, run the following command to remove those permissions so the mode is 0755 or more restrictive:

```
# chmod 0755 <directory>
```

For each PATH entry that does not exist on the system, either remove it from root's PATH or run the following command to create it with root ownership and an appropriate mode:

```
# mkdir -p <directory> && chown root <directory> && chmod 0755 <directory>
```

After making these changes, re-run the audit to confirm root's PATH contains only existing, root-owned directories that are not writable by group or other, with no empty, trailing, or . entries.

References:

1. GRID: NIX-00000256
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			

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`) determines the permissions applied to newly created files and directories. On Linux the default permissions are `0777` (rwxrwxrwx) for a new directory and `0666` (rw-rw-rw-) for a new file; the `umask` restricts (masks) these defaults bitwise - bits set in the `umask` are cleared in the resulting mode.

A `umask` can be expressed as an octal or a symbolic value:

- Octal (numeric) value - three or four digits, for example `umask 0027` or `umask 027`. When four digits are given the leading digit is ignored; the remaining three set the masking for user, group, and other respectively.
- Symbolic value - a comma-separated list of the permissions to keep for user (`u`), group (`g`), and other (`o`). For example, `umask u=rwx,g=rx,o=` is equivalent to the octal `umask 027`, producing a new directory of `drwxr-x---` and a new file of `rw-r-----`.

root's effective `umask` is the value in force after every source that sets it (the system defaults and root's own shell initialization files) has been applied in order of precedence.

Rationale:

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 display the `root` user's effective login `umask` - the value in force after the system defaults and root's shell initialization files have been applied:

```
# su - root -c 'umask'
```

The value returned should be `0027` (or the equivalent `027`) or more restrictive - that is, every permission bit denied by `0027` is also denied by the effective `umask`. A returned value of `0027`, `0077`, or any value at least as restrictive is compliant; a less restrictive value such as `0022` is not.

Remediation:

Edit /root/.bash_profile and /root/.bashrc and set root's umask to 0027 or more restrictive by adding or modifying the umask line to read:

```
umask 027
```

After making the change, verify that root's effective login umask resolves to 0027 or more restrictive, since a umask set in a system-wide location (for example /etc/profile or a file under /etc/profile.d/) can override the value set in root's shell initialization files.

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

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.	●	●	●

5.4.2.7 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:

A system account that retains a valid login shell provides an interactive foothold for anyone who obtains its credentials. Removing the login shell from accounts that do not need interactive access eliminates that foothold.

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" && $NF != "false")
{print}' /etc/shells | sed -rn '/^\//{s,/,\\\\/,g;p}' | paste -s -d '|' -
))$"
   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:

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

{
   l_valid_shells="^($(awk -F\/ '($NF != "nologin" && $NF != "false")
{print}' /etc/shells | sed -rn '/^\//{s,/,\\\\/,g;p}' | paste -s -d '|' -
))$"
   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`, `halt`, `sync`, `shutdown`, and `nfsnobody` users are exempted from requiring a non-login shell.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	4.7 Manage Default Accounts on Enterprise Assets and Software Manage default accounts on enterprise assets and software, such as root, administrator, and other pre-configured vendor accounts. Example implementations can include: disabling default accounts or making them unusable.	●	●	●
v7	16.8 Disable Any Unassociated Accounts Disable any account that cannot be associated with a business process or business owner.	●	●	●

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. Additionally, an administrator may add special accounts that are not intended for interactive use.

Rationale:

An account without a valid login shell can still be used for authentication unless it is locked. Locking such accounts closes that residual authentication path.

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" && $NF != "false")
{print}' /etc/shells | sed -rn '/^\//{s,/,\\\\/,g;p}' | paste -s -d '|' -
))$"
   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:

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

{
   l_valid_shells="^($(awk -F\/ '($NF != "nologin" && $NF != "false")
{print}' /etc/shells | sed -rn '/^\//{s,/,\\\\/,g;p}' | paste -s -d '|' -
))$"
   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	4.7 Manage Default Accounts on Enterprise Assets and Software Manage default accounts on enterprise assets and software, such as root, administrator, and other pre-configured vendor accounts. Example implementations can include: disabling default accounts or making them unusable.	●	●	●
v7	16.8 Disable Any Unassociated Accounts Disable any account that cannot be associated with a business process or business owner.	●	●	●

5.4.3 Configure user default environment

This subsection covers settings that shape a user's default runtime environment on the system: the login shells permitted in /etc/shells, the interactive shell idle timeout (TMOUT), and the default file-creation mode mask (umask).

Several of these are governed by the system-wide shell-initialization files, which a login shell sources in order after the PAM session is established: /etc/profile, then the files under /etc/profile.d/ (including /etc/profile.d/sh.local, which /etc/profile sources but which does not match a *.sh glob), then /etc/bashrc. A setting applied later in this sequence overrides an earlier one. These recommendations address the bash and sh shell-initialization files; if other login shells are supported on the system, their equivalent configuration should be reviewed as well, as it is not covered here.

The default user umask is addressed by two complementary recommendations: one verifies the UMASK baseline in /etc/login.defs (the value pam_umask applies before the shell starts, and the default for sessions that never source the shell-initialization files), and the other verifies that no shell-initialization file relaxes that baseline below the required floor. Together they ensure the effective umask a user receives -- interactive or not -- is 0027 or more restrictive. A default umask is a baseline rather than a hard limit: a user can apply a more permissive umask for their own session (for example through ~/.bashrc, the umask command, or chmod on individual files), so any other mechanism that sets a umask should be reviewed as well.

These are host-level account-environment settings and apply across all deployment types, including cloud/CSP images.

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 `/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)

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			

5.4.3.2 Ensure default user shell timeout is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

TMOUT is a shell environment variable that sets an interactive shell's idle timeout, in seconds: after `TMOUT` seconds with no input, the shell exits. `TMOUT=<n>` sets the timeout to n seconds (a value of `0` disables it); marking the variable `readonly` prevents a user from overriding it during the session, and `export`ing it applies it to the user's shell environment.

Rationale:

Setting a timeout value reduces the window of opportunity for unauthorized user access to another user's shell session that has been left unattended. It also ends the inactive session and releases the resources associated with that session.

Audit:

Run the following script to verify that TMOUT is set to a value greater than 0 and no larger than 900 seconds, and is marked readonly and exported, in each shell initialization file that configures it:

```
#!/usr/bin/env bash
{
   l_max=900
   a_files=()
   while IFS= read -r l_file; do a_files+=("$l_file"); done < <(grep -Pls --
'^([^#\n\r]+)?\bTMOUT\b' \
      /etc/*bashrc /etc/profile /etc/profile.d/*.sh /etc/profile.d/sh.local \
      /usr/etc/profile.d/*.sh /usr/etc/profile.d/sh.local 2>/dev/null)
   l_set=0
   for l_file in "${a_files[@]}"; do
      l_val="$(grep -Po -- '^([^#\n\r]+)?\bTMOUT=\d+\b' "$l_file" | awk -F=
'{v=$2} END{if (NR>0) print v}')"
      [ -z "$l_val" ] && continue
      l_set=1
      l_ro="$(grep -P -- '^\h*(typeset\h\-
xr\hTMOUT=\d+|([^#\n\r]+)?\breadonly\h+TMOUT\b)' "$l_file")"
      l_exp="$(grep -P -- '^\h*(typeset\h\-
xr\hTMOUT=\d+|([^#\n\r]+)?\bexport\b([^#\n\r]+\b)?TMOUT\b)' "$l_file")"
      [ "$l_val" -gt 0 ] && [ "$l_val" -le "$l_max" ] || echo "$l_file:
TMOUT=$l_val is out of range (must be >0 and <=$l_max)"
      [ -n "$l_ro" ]  || echo "$l_file: TMOUT is not marked readonly"
      [ -n "$l_exp" ] || echo "$l_file: TMOUT is not exported"
   done
   [ "$l_set" -eq 0 ] && echo "TMOUT is not set in any shell initialization
file"
}
```

Nothing should be returned.

Remediation:

Configure TMOUT in a single file. Because best practice marks TMOUT readonly, a second assignment from a later-sourced file produces a "TMOUT: readonly variable" error at every login, so it must be set exactly once.

1. Run the following command to locate any existing TMOUT assignments in the shell initialization files:

```
# grep -PHrs -- '\bTMOUT\b' /etc/profile /etc/profile.d/ /etc/bashrc
/etc/*bashrc
```

2. In each file returned in step one, remove or comment out the TMOUT lines so TMOUT is no longer set in more than one location.
3. Create a single drop-in file in /etc/profile.d/ that sets, marks readonly, and exports TMOUT in one statement, using a value greater than 0 and no greater than 900 seconds in line with local site policy.

Example:

```
# printf '%s\n' "# Set a 900 second interactive shell timeout" "typeset -xr
TMOUT=900" > /etc/profile.d/50-tmout.sh
```

References:

1. GRID: NIX-00000260

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.	●	●	●

5.4.3.3 Ensure default user umask is configured in /etc/login.defs (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

`/etc/login.defs` sets system-wide account defaults. Its `UMASK` value is the umask baseline that `pam_umask` applies to a login session before any shell-initialization file runs, and that `useradd` uses when creating a home directory. It governs the umask for sessions that do not source the interactive shell-initialization files.

Rationale:

If the `/etc/login.defs` `UMASK` is missing or less restrictive than `0027`, sessions that never run the shell-initialization files -- and files created during account setup -- receive permissions readable or writable by other users. A restrictive baseline here ensures a safe default even where shell-level umask settings do not apply.

Audit:

Run the following command to verify the `UMASK` value in `/etc/login.defs` is `0027` or more restrictive:

```
# grep -Pi -- '^\h*UMASK\h+\H+' /etc/login.defs
```

A `UMASK` value of `0027` (or `027`) or more restrictive with `UMASK` being uppercase should be returned.

Example output:

```
UMASK 0027
```

Remediation:

Edit `/etc/login.defs` and set the `UMASK` line to `0027` or more restrictive:

```
UMASK 027
```

References:

1. GRID: NIX-00000389
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)

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.	●	●	●

5.4.3.4 Ensure default user umask is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The user file-creation mode mask (`umask`) determines the permissions withheld from newly created files and directories. A user's effective `umask` is the value in force after login; the shell-initialization files `/etc/profile`, files under `/etc/profile.d/`, and `/etc/bashrc` can set it and can override the system baseline. This recommendation covers those shell-initialization sources; the baseline value in `/etc/login.defs` is covered by the separate recommendation Ensure default user umask is configured in /etc/login.defs.

Rationale:

A `umask` less restrictive than `0027` causes new files and directories to be readable or writable by users other than the owner. A shell-initialization file that sets a permissive `umask` overrides a restrictive baseline for every interactive login, extending those group- or other-accessible defaults to every file created in that session.

Audit:

Run the following script to verify that none of the shell-initialization sources -- /etc/profile, the files ending in .sh under /etc/profile.d/, /etc/profile.d/sh.local, /etc/bashrc, and /etc/bash.bashrc -- set a umask less restrictive than 0027. Both octal and symbolic umask statements are evaluated:

```
#!/usr/bin/env bash
{
   l_floor=$(( 8#0027 ))

   # Convert a umask argument (octal "0027" or symbolic "u=rwx,g=rx,o=") to
its
   # decimal deny-bits value; prints "INVALID" for a form that cannot be
   # evaluated statically (relative +/- operators or unknown tokens).
   f_umask_to_dec() {
      local l_arg="$1" l_u=0 l_g=0 l_o=0 l_clause l_who l_perms l_bits l_deny
      if [[ "$l_arg" =~ ^[0-7]{3,4}$ ]]; then
         printf '%d' "$(( 8#$l_arg ))"; return 0
      fi
      local IFS=','
      local -a a_clauses=( $l_arg )
      for l_clause in "${a_clauses[@]}"; do
         if [[ "$l_clause" =~ ^([augo]+)=([rwx]*)$ ]]; then
            l_who="${BASH_REMATCH[1]}"; l_perms="${BASH_REMATCH[2]}";
l_bits=0
            [[ "$l_perms" == *r* ]] && l_bits=$(( l_bits | 4 ))
            [[ "$l_perms" == *w* ]] && l_bits=$(( l_bits | 2 ))
            [[ "$l_perms" == *x* ]] && l_bits=$(( l_bits | 1 ))
            l_deny=$(( 7 - l_bits ))
            [[ "$l_who" == *a* ]] && { l_u=$l_deny; l_g=$l_deny; l_o=$l_deny;
}
            [[ "$l_who" == *u* ]] && l_u=$l_deny
            [[ "$l_who" == *g* ]] && l_g=$l_deny
            [[ "$l_who" == *o* ]] && l_o=$l_deny
         else
            printf 'INVALID'; return 0
         fi
      done
      printf '%d' "$(( l_u * 64 + l_g * 8 + l_o ))"
   }

   a_files=(/etc/profile /etc/bashrc /etc/bash.bashrc
/etc/profile.d/sh.local)
   while IFS= read -r -d '' l_file; do
      a_files+=("$l_file")
   done < <(find /etc/profile.d/ -type f -name '*.sh' -print0 2>/dev/null)

   for l_file in "${a_files[@]}"; do
      [ -r "$l_file" ] || continue
      grep -Psi -- '^\h*umask\h+\S' "$l_file" 2>/dev/null | while IFS= read -
r l_line; do
         l_arg="$(printf '%s' "$l_line" | grep -Poi --
'^\h*umask\h+\K[^[:space:]#]+')"
         [ -n "$l_arg" ] || continue
         l_dec="$(f_umask_to_dec "$l_arg")"
         { [ "$l_dec" = "INVALID" ] || [ "$(( l_dec & l_floor ))" -ne
"$l_floor" ]; } && printf '%s: %s\n' "$l_file" "$l_line"
      done
   done
}
```

Nothing should be returned.

Remediation:

Identify any shell-initialization umask statement less restrictive than 0027, and tighten or remove it. When no shell-init source sets a umask, the /etc/login.defs baseline (a separate recommendation) governs the default.

1. Locate the umask statements in the shell-initialization files:

```
# grep -PHrsi -- '^\h*umask\b' /etc/profile /etc/profile.d/ /etc/bashrc
/etc/bash.bashrc
```

2. For each statement returned that is less restrictive than 0027, edit the file to set a compliant value:

```
umask 027
```

or remove the statement so the /etc/login.defs baseline governs the default.

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

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 Logging and Auditing

The items in this section describe how to configure logging, log monitoring, and auditing, using tools included in most distributions.

It is recommended that `rsyslog` be used for logging (with `logwatch` providing summarization) and `auditd` be used for auditing (with `aureport` providing summarization) to automatically monitor logs for intrusion attempts and other suspicious system behavior.

In addition to the local log files created by the steps in this section, it is also recommended that sites collect copies of their system logs on a secure, centralized log server via an encrypted connection. Not only does centralized logging help sites correlate events that may be occurring on multiple systems, but having a second copy of the system log information may be critical after a system compromise where the attacker has modified the local log files on the affected system(s). If a log correlation system is deployed, configure it to process the logs described in this section.

Because it is often necessary to correlate log information from many different systems (particularly after a security incident) it is recommended that the time be synchronized among systems and devices connected to the local network. The standard Internet protocol for time synchronization is the Network Time Protocol (NTP), which is supported by most network-ready devices. Reference <http://chrony.tuxfamily.org/> manual page for more information on configuring chrony.

It is important that all logs described in this section be monitored on a regular basis and correlated to determine trends. A seemingly innocuous entry in one log could be more significant when compared to an entry in another log.

Note on log file permissions: There really isn't a "one size fits all" solution to the permissions on log files. Many sites utilize group permissions so that administrators who are in a defined security group, such as "wheel" do not have to elevate privileges to root in order to read log files. Also, if a third party log aggregation tool is used, it may need to have group permissions to read the log files, which is preferable to having it run setuid to root. Therefore, there are two remediation and audit steps for log file permissions. One is for systems that do not have a secured group method implemented that only permits root to read the log files (`root:root 600`). The other is for sites that do have such a setup and are designated as `root:securegrp 640` where `securegrp` is the defined security group (in some cases `wheel`).

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 principles 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 covers system logging on a default installation, where systemd-journald and rsyslog work together: systemd-journald is always running and is the first collector of log data, and rsyslog is installed and enabled by default, reads the messages that journald collects, and writes the persistent text logs under /var/log. This guidance focuses on rsyslog as the default persistent logging service and configures journald in its role as the collector that feeds rsyslog. It also covers restricting access to the log files under /var/log/ so their contents cannot be read or altered by unauthorized users.

Configuring journald as a standalone persistent log store in place of rsyslog is out of scope: these recommendations focus on the default rsyslog path and do not cover running journald as the sole logging system.

If your organization makes use of an enterprise-wide logging system completely outside of rsyslog and journald, the following recommendations do not directly apply; however, the principles 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 relevant recommendations to determine exactly what applies.

What is not covered

- Enterprise logging systems not utilizing rsyslog or journald. As logging is very situational and dependent on the local environment, not everything can be covered here.
- Transport layer security should be applied to all remote logging functionality. Both rsyslog and journald support 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

This section provides recommendations for configuring the systemd journal service, systemd-journald. The recommendations ensure the journald service is active, that it is configured to forward events to rsyslog, that its journal files are rotated and their access is restricted, and that it is not used as an unintended remote log receiver.

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 received from a variety of sources such as:

- Classic RFC3164 BSD syslog via the /dev/log socket
- STDOUT/STDERR of programs via StandardOutput=journal and 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

On this platform systemd-journald is a static, socket-activated unit that is always present and is started by systemd rather than enabled on boot. It serves as the initial collector of log events and feeds rsyslog, the default persistent logging service (see the System Logging section overview). By default journald stores its journals in volatile memory, and rsyslog provides persistent on-disk retention under /var/log; the two coexist rather than replace one another. Configuring journald as a standalone persistent log store or as a remote-logging mechanism in place of rsyslog is out of scope.

journald is configured in /etc/systemd/journald.conf and in drop-in files ending in .conf under /etc/systemd/journald.conf.d/. Drop-in files are the recommended method for local configuration; administrator drop-ins use the 60-90 filename range, and where the same setting appears more than once the highest-precedence (last-sorted) entry wins. The effective configuration can be reviewed with systemd-analyze cat-config systemd/journald.conf.

Any changes made to the systemd-journald configuration will require a restart of systemd-journald.

6.1.1.1 Ensure journald service is active (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

systemd-journald is the systemd journaling service that collects and stores system log data. It is a static, socket-activated unit that is always present as the initial collector of log events on a systemd-based system; it is started by systemd rather than enabled to start on boot.

Rationale:

systemd-journald must be running for the system to capture log events. If the service is not active, log events are not collected, removing the local record needed to detect, investigate, and respond to security-relevant activity.

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 and start systemd-journald.service:

```
# systemctl unmask systemd-journald.service
# systemctl start 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

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.		●	●

6.1.1.2 Ensure systemd-journal-remote is not in use (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The systemd-journal-remote package provides the ability to receive journal log entries from remote hosts over the network, which makes the receiving host a log server. The same package supplies both the sending and the receiving functions; the receiving side is handled by two units, systemd-journal-remote.socket and systemd-journal-remote.service.

Rationale:

If a client is configured to also receive log data, thus turning it into a log server, the client system is acting outside its operational boundary. Disabling the systemd-journal-remote receiver is collector hygiene that applies regardless of whether rsyslog or journald is the primary log path.

Audit:

Run the following command to verify the systemd-journal-remote receiver is not in use. It reports any of its units that are enabled (including enabled-runtime) or active:

```
# systemctl show -p Id -p UnitFileState -p ActiveState systemd-journal-
remote.socket systemd-journal-remote.service | awk -v RS='' '{ id=""; u=0;
for(i=1;i<=NF;i++){ if($i ~ /^Id=/) id=substr($i,4); if($i ~
/^UnitFileState=(enabled|enabled-runtime)$/) u=1; if($i ~
/^ActiveState=active$/) u=1 } if(u) print id " is in use" }'
```

Nothing should be returned.

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.		●	●

6.1.1.3 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:

systemd-journald is always present and acts as the initial collector of system log events. Forwarding those events to rsyslog provides a single, consistent path for log processing and export, including reliable delivery to a centralized log server for retention and correlation. When forwarding is not enabled, events collected by journald are not passed to rsyslog for further processing.

Impact:

With forwarding enabled, events collected by systemd-journald are duplicated into rsyslog in addition to being retained by journald. Ensure rsyslog log storage and rotation are sized to accommodate the additional log volume this forwarding generates.

Audit:

Run the following command to verify that `ForwardToSyslog` is effectively set to yes, using the highest-precedence active setting across the main configuration file and any drop-ins:

```
# systemd-analyze cat-config systemd/journald.conf | grep -Pi --
'^\h*ForwardToSyslog\h*=' | tail -n 1
```

Verify the output:

```
ForwardToSyslog=yes
```

Remediation:

Ensure `ForwardToSyslog` is set to `yes` and that no other entry overrides it.

1. Run the following command to locate every file that sets `ForwardToSyslog`:

```
# grep -PHrs -- '^\h*ForwardToSyslog\h*=' /etc/systemd/journald.conf
/etc/systemd/journald.conf.d/
```

2. In each file returned, remove or comment out the `ForwardToSyslog` line so the setting is defined in only one place.
3. Set `ForwardToSyslog=yes` in a drop-in file in `/etc/systemd/journald.conf.d/` ending in `.conf`.

Example:

```
# [ ! -d /etc/systemd/journald.conf.d ] && mkdir -p
/etc/systemd/journald.conf.d
# printf '%s\n' "[Journal]" "ForwardToSyslog=yes" >>
/etc/systemd/journald.conf.d/60-journald.conf
```

Run the following command to apply the change to the running service:

```
# systemctl reload-or-restart systemd-journald.service
```

References:

1. GRID: NIX-00000265
2. NIST SP 800-53 :: AU-6
3. SYSTEMD-JOURNALD.SERVICE(8)
4. JOURNALD.CONF(5)
5. systemd journald.conf(5): https://www.freedesktop.org/software/systemd/man/latest/journald.conf.html

Additional Information:

Events generated very early in boot, before the forwarding path is available, are retained by journald and may not be forwarded.

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.		●	●

6.1.1.4 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. Journal rotation is configured in /etc/systemd/journald.conf and in drop-in files under /etc/systemd/journald.conf.d/.

Rationale:

By keeping the log files smaller and more manageable, a system administrator can easily archive these files to another system and spend less time looking through inordinately large log files.

Audit:

Review /etc/systemd/journald.conf and the files ending in .conf under /etc/systemd/journald.conf.d/, and verify that journal rotation is configured according to site policy. The following command lists the effective (highest-precedence) rotation settings:

```
# systemd-analyze cat-config systemd/journald.conf | tac | grep -Psi --
'^\h*(SystemMaxUse|SystemKeepFree|SystemMaxFileSize|RuntimeMaxUse|RuntimeKeep
Free|RuntimeMaxFileSize|MaxFileSec)\h*='
```

One or more of the following should be set according to site policy:

```
SystemMaxUse=
SystemKeepFree=
SystemMaxFileSize=
RuntimeMaxUse=
RuntimeKeepFree=
RuntimeMaxFileSize=
MaxFileSec=
```

Remediation:

Set the journal rotation parameters in a drop-in file in /etc/systemd/journald.conf.d/ ending in .conf. Set the parameters in the [Journal] section, adjusting the values to site policy, and ensure no other file overrides them.

Example:

```
[Journal]
SystemMaxUse=1G
SystemKeepFree=500M
SystemMaxFileSize=100M
RuntimeMaxUse=200M
RuntimeKeepFree=50M
RuntimeMaxFileSize=50M
MaxFileSec=1month
```

Note: These parameters have specific edge cases and interactions; review journald.conf(5) for the exact semantics.

References:

1. GRID: NIX-00000267
2. NIST SP 800-53 Rev. 5: AU-2, AU-7, AU-12
3. systemd journald.conf(5): https://www.freedesktop.org/software/systemd/man/latest/journald.conf.html

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.		●	●

6.1.1.5 Ensure journald log file access is configured (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Journald creates log files that do not already exist on the system. The permissions applied to these newly created files are governed by systemd-tmpfiles directives in /usr/lib/tmpfiles.d/systemd.conf, which an administrator may override in /etc/tmpfiles.d/systemd.conf.

Rationale:

Journal files can contain sensitive information such as usernames, IP addresses, and details of system and user activity. Restricting access to these files to their owner and group prevents unauthorized users from reading that information or tampering with the local log record.

Audit:

Journald applies the modes defined by systemd-tmpfiles directives in /usr/lib/tmpfiles.d/systemd.conf (or an override in /etc/tmpfiles.d/systemd.conf, which takes precedence) to the log files it creates. Existing files keep their current mode until recreated, so review both the directive and the files on disk.

1. Run the following command to review the mode field (the third column) of the journald log file directives:

```
# systemd-tmpfiles --cat-config | grep -P -- '/var/log/journal/.*\.journal'
```

Each returned mode should be 0640 or more restrictive

2. Run the following command to list any journald log files on disk with permissions less restrictive than `0640`:

```
# find /var/log/journal /run/log/journal -type f -name '*.journal*' -perm
/0137
```

Nothing should be returned

Remediation:

Set the journald log file mode to `0640` (or a more restrictive value per site policy), then correct any existing files.

1. Run the following command to create a local override of the tmpfiles configuration, if it does not already exist, so the change survives package updates:

```
# [ -f /etc/tmpfiles.d/systemd.conf ] || cp /usr/lib/tmpfiles.d/systemd.conf
/etc/tmpfiles.d/systemd.conf
```

2. Edit /etc/tmpfiles.d/systemd.conf and set the mode field (the third column) of each journald log file directive to 0640 or more restrictive.

Run the following command to apply the updated tmpfiles configuration:

```
# systemd-tmpfiles --create
```

Run the following command to correct the mode of any existing journald log files that are looser than 0640 (this removes only the excess permission bits and does not loosen files that are already more restrictive):

```
# find /var/log/journal /run/log/journal -type f -name '*.journal*' -perm
/0137 -exec chmod u-x,g-wx,o-rwx {} +
```

References:

1. GRID: NIX-00000266
2. NIST SP 800-53 :: AU-9, AC-3, AC-6
3. systemd tmpfiles.d(5): https://www.freedesktop.org/software/systemd/man/latest/tmpfiles.d.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.	●	●	●

6.1.2 Configure rsyslog

This section provides recommendations for configuring the rsyslog logging service. The recommendations ensure that rsyslog is installed and active, that the log files it creates and its own configuration and spool files are appropriately protected, that the host does not act as a remote log server, that local logging and - where used - encrypted forwarding to a central log host are configured, and that logs are rotated to bound disk use.

rsyslog is installed and enabled by default and is the recommended, most commonly deployed persistent logging service on this platform. It reads the messages that systemd-journald collects and writes them to the text log files under /var/log. This guidance treats rsyslog as the default persistent logging path and remediates toward a secure, best-practice rsyslog configuration; systemd-journald remains present as the collector that feeds it and is addressed in the Configure journald subsection. Running systemd-journald as a standalone persistent store in place of rsyslog is out of scope. Where rsyslog is not installed or otherwise not in use - for example, a different logging system is in use and approved by local site policy - the recommendations in this section are Not Applicable and the section may be skipped.

This subsection follows the rsyslog lifecycle:

- Foundation - install the rsyslog package and enable and activate the service.
- File-permission posture - constrain the mode of the log files rsyslog CREATES (the log file creation mode / umask), and separately restrict ownership and access to rsyslog's OWN configuration files (/etc/rsyslog.conf and the .conf drop-ins under /etc/rsyslog.d/), the /etc/rsyslog.d/ directory, and the work (queue-spool) directory (default /var/lib/rsyslog). These are distinct concerns: the creation-mode recommendation governs the permissions rsyslog applies to the log files it WRITES, whereas the access recommendations govern who may read or modify rsyslog's own configuration and spool files. These configuration and spool objects are static and administrator-managed - rsyslog does not recreate them or reset their permissions at runtime - so the hardened modes are durable once applied.
- Logging behavior - ensure the host does not act as a log server (does not receive remote messages), that local logging rules are configured, and, where used, that logs are forwarded to a central log host.
- Encrypted transport - where logs are forwarded over the network, secure that transport with TLS: install a TLS network stream driver, require the forwarding action to use TLS, configure a trusted CA certificate so the sender validates the receiver, and - where the receiver requires client authentication - configure a client certificate so the sender authenticates to the receiver (mutual TLS). These apply only where remote forwarding over TLS is used; see the rsyslog documentation for driver and certificate implementation details.
- Retention - rotate log files with logrotate to bound disk use.

rsyslog configuration may be written in three formats: the basic (sysklogd) format; the advanced (RainerScript; global(...) / action(...)) format, which is the current best practice; and the obsolete legacy $-directive format, which is deprecated and should be migrated to the advanced format. Audit and remediation steps in this subsection accept the advanced and legacy forms where both exist; write new configuration in the advanced format.

6.1.2.1 Ensure rsyslog is installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

rsyslog is a reliable, feature-rich system logging service. The `rsyslog` package provides the rsyslogd daemon, which collects log messages from the kernel, system services, and applications (including those gathered by the systemd journal) and writes them to persistent on-disk log files under /var/log, and can filter, rate-limit, and forward them to a central log host.

Rationale:

Persistent, reliable logging is fundamental to security operations: it preserves the record of system and security events needed for accountability, threat detection, incident investigation, and forensic analysis. The systemd journal stores its logs in volatile memory by default, so without a persistent logging service those records can be lost across a reboot or under memory pressure. Installing rsyslog establishes durable on-disk log retention and the foundation for later centralization and protection of that log data.

Impact:

None. rsyslog is installed and enabled by default on supported platforms; where the package must be installed, the operational impact is negligible.

Audit:

Run the following command to verify the `rsyslog` package is installed:

```
# rpm -q rsyslog
```

Example output:

```
rsyslog-8.2402.0-1
```

A package name and version should be returned.

Remediation:

Run the following command to install `rsyslog`:

```
# dnf install rsyslog
```

Default Value:

Installed

References:

1. GRID: NIX-00000276
2. NIST SP 800-53 :: CM-6 b
3. NIST SP 800-53A :: CM-6.1 (iv)
4. RHEL 8 STIG GROUP ID: V-230477
5. RHEL 8 STIG RULE ID: SV-230477r627750

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.		●	●

6.1.2.2 Ensure rsyslog service is enabled and active (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The rsyslog service (rsyslogd) runs the logging daemon. It must be enabled so it starts automatically at boot and active so it is running in the current session; an enabled-and-active service ensures log collection begins at boot and continues without manual intervention.

Rationale:

If the rsyslog service is not running and enabled at boot, the system stops capturing and persisting log messages, leaving gaps in the record of system and security events that accountability, detection, and investigation depend on.

Impact:

None. rsyslog is enabled and active by default on supported platforms; enabling and starting an installed service has negligible operational impact.

Audit:

1. Run the following command to verify `rsyslog.service` is enabled:

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

Example output:

```
enabled
```

2. Run the following command to verify `rsyslog.service` is active:

```
# systemctl is-active rsyslog.service
```

Example output:

```
active
```

Remediation:

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. 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.		●	●

6.1.2.3 Ensure rsyslog log file creation mode is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

rsyslog sets the permissions on the log files it creates. The file creation mode is controlled by the `fileCreateMode` parameter - the modern `global(fileCreateMode="...")` / omfile-action form or the legacy `$FileCreateMode` directive - and is further constrained by rsyslog's umask (`global(umask=...)` or the legacy `$umask`). The value is a four-digit octal mode; if no creation mode is set, rsyslog defaults to 0644. This recommendation governs that creation mode so newly created log files are not world-readable.

Rationale:

Log files routinely contain sensitive event data - usernames, source addresses, service and command detail, and occasionally secrets logged in error. If rsyslog creates log files readable by group or other, any local user can read that data. Constraining the creation mode to 0640 or more restrictive ensures newly created logs are readable only by root (and the log group).

Impact:

Restricting the creation mode removes read access for accounts other than root (and the root/log group); rsyslog runs as root and is unaffected. Because rsyslog's default creation mode is 0644 (world-readable), applying this recommendation changes the permissions of newly created log files on a default system.

Audit:

Run the following script to assemble the effective rsyslog configuration and verify the log file creation mode is 0640 or more restrictive. The check matches both the modern advanced `fileCreateMode=` parameter (for example `global(fileCreateMode="0640")`) and the legacy `$FileCreateMode` directive; a match on either satisfies this recommendation, though the legacy `$`-directive form is deprecated and should be migrated to the advanced format:

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

{
   a_conf=(/etc/rsyslog.conf /etc/rsyslog.d/*.conf)
   while IFS= read -r l_inc; do
      for l_f in $l_inc; do
         [ -d "$l_f" ] && a_conf+=("$l_f"/*.conf) || a_conf+=("$l_f")
      done
   done < <(grep -Psoih --
'^\h*(?:\$IncludeConfig\h+|include\h*\(\h*file\h*=\h*"?)\K[^")\h]+'
"${a_conf[@]}" 2>/dev/null)
   grep -Psi -- '^\h*(\$FileCreateMode\h+|.*\bfileCreateMode\h*=\h*"?)0[0-
7]{3}' "${a_conf[@]}" 2>/dev/null
}
```

Example output (the first line is the modern advanced form, the second the equivalent legacy directive; either satisfies):

```
global(fileCreateMode="0640")
$FileCreateMode 0640
```

A value of `0640` or more restrictive (for example `0640`, `0600`, `0400`) should be configured. Where site policy dictates a different mode, follow that policy.

Note: A `umask` setting (`$umask` or `global(umask=)`) is a separate, coarser control that masks permission bits and does not by itself satisfy this recommendation; the file creation mode must be set via `fileCreateMode` or `$FileCreateMode`.

Remediation:

1. Run the following command to list any existing file-creation-mode directives:

```
# grep -PHrs -- '^\h*(\$FileCreateMode\b|[^#]*\bfileCreateMode\h*=)'
/etc/rsyslog.conf /etc/rsyslog.d/
```

Remove or correct any `$FileCreateMode` / `fileCreateMode=` line more permissive than `0640` so the mode is defined once

2. Create a drop-in file (for example `/etc/rsyslog.d/20-file-create-mode.conf`) with the following content, setting the mode to `0640` or more restrictive:

```
global(fileCreateMode="0640")
```

3. Run the following command to restart `rsyslog`:

```
# systemctl restart rsyslog
```

The legacy directive `$FileCreateMode 0640` is equivalent and remains supported where the modern syntax is not used.

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.		●	●

6.1.2.4 Ensure access to rsyslog configuration files is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The main rsyslog configuration file /etc/rsyslog.conf and the drop-in configuration files ending in .conf under /etc/rsyslog.d/ control how rsyslog collects, filters, stores, and forwards system log data. These files should be owned by root, group-owned by root, and no more permissive than mode 0640, so that they can be modified only by root and are not readable by any account other than root.

Rationale:

The rsyslog configuration governs which events are recorded and where they are sent. If a non-privileged account can modify these files, it could disable logging of critical events, redirect logs away from a central collector, or alter filters to hide activity, undermining the integrity and availability of the audit trail. Restricting access to root protects the logging configuration from unauthorized tampering and disclosure.

Impact:

The recommended mode 0640 removes read access for accounts other than root and the root group; rsyslog runs as root and is unaffected. The rsyslog package ships these files world-readable (mode 0644) by default, so applying this recommendation changes their permissions on a default system. This is consistent with the more restrictive 0640 convention already applied to the audit configuration files in this benchmark.

Audit:

Verify that the rsyslog configuration files are owned by root, group-owned by root, and mode 0640 or more restrictive.

1. Run the following command to list the main configuration file if it is not owned by root, not group-owned by root, or less restrictive than mode 0640:

```
# find /etc/rsyslog.conf -not -type l \( ! -user root -o ! -group root -o -
perm /0137 \) 2>/dev/null
```

Nothing should be returned.

2. Run the following command to list any .conf drop-in under /etc/rsyslog.d/ that is not owned by root, not group-owned by root, or less restrictive than mode 0640:

```
# find /etc/rsyslog.d/ -type f -name '*.conf' \( ! -user root -o ! -group
root -o -perm /0137 \) 2>/dev/null
```

Nothing should be returned.

Note: If rsyslog is not installed, /etc/rsyslog.conf and /etc/rsyslog.d/ may not exist; the commands return nothing and this recommendation is Not Applicable.

Remediation:

Run the following commands to set ownership to root:root and remove group write, all access for other accounts, and all execute bits from the rsyslog configuration files, resulting in mode 0640 or more restrictive:

```
# chown root:root /etc/rsyslog.conf
# chmod u-x,g-wx,o-rwx /etc/rsyslog.conf
# find /etc/rsyslog.d/ -type f -name '*.conf' -exec chown root:root {} +
# find /etc/rsyslog.d/ -type f -name '*.conf' -exec chmod u-x,g-wx,o-rwx {} +
```

Default Value:

/etc/rsyslog.conf is mode 0644, owner root, group root; drop-in .conf files under /etc/rsyslog.d/ are mode 0644, owner root, group root (EL9 rsyslog package default). Confirm the stock modes on a representative host with rpm -qlv rsyslog and stat -c '%a %U %G' /etc/rsyslog.conf before finalizing.

References:

1. GRID: NIX-00000426
2. https://www.rsyslog.com/doc/configuration/index.html

Additional Information:

The target mode 0640 (owner read and write, group read, no access for other accounts) matches the convention applied to the audit configuration files in this benchmark and is more restrictive than the rsyslog package default of 0644. These are static, administrator-managed configuration files: rsyslog does not recreate them or reset their permissions at runtime, so the hardened mode is durable once applied.

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.	●	●	●

6.1.2.5 Ensure access to the /etc/rsyslog.d directory is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The /etc/rsyslog.d/ directory holds drop-in rsyslog configuration files ending in .conf. The directory should be owned by root, group-owned by root, and not writable by any account other than root so that unauthorized users cannot add, remove, or replace rsyslog configuration fragments.

Rationale:

Because rsyslog reads every .conf file in /etc/rsyslog.d/, an account able to write to this directory could introduce a new configuration fragment that disables logging, changes filters, or redirects log forwarding, without altering any existing file. Restricting write access to the directory to root prevents unauthorized modification of the effective logging configuration.

Impact:

None. The recommended ownership and permissions match the rsyslog package default (0755 root:root), so applying this recommendation does not change default system behavior.

Audit:

Verify that the /etc/rsyslog.d/ directory is owned by root, group-owned by root, and mode 0755 or more restrictive (that is, not group- or world-writable).

Run the following command to list the directory if it is not owned by root, not group-owned by root, or less restrictive than mode 0755:

```
# find /etc/rsyslog.d/ -maxdepth 0 -type d \( ! -user root -o ! -group root -
o -perm /0022 \) 2>/dev/null
```

Nothing should be returned.

Note: If rsyslog is not installed, /etc/rsyslog.d/ may not exist; the command returns nothing and this recommendation is Not Applicable.

Remediation:

Run the following commands to set ownership to root:root and remove group and other write from the directory:

```
# chown root:root /etc/rsyslog.d
# chmod g-w,o-w /etc/rsyslog.d
```

Default Value:

/etc/rsyslog.d/ is mode 0755, owner root, group root (EL9 rsyslog package default).

References:

1. GRID: NIX-00000427
2. https://www.rsyslog.com/doc/configuration/index.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.	●	●	●

6.1.2.6 Ensure access to the rsyslog work directory is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

rsyslog's work directory, set by the global $WorkDirectory / global(workDirectory=...) directive and defaulting to /var/lib/rsyslog, stores disk-assisted queue spool files. These files can hold buffered log data written to disk when in-memory queues are exhausted or a forwarding action is retrying. The work directory should be owned by root, group-owned by root, and accessible only to root.

Rationale:

Spooled queue files under the rsyslog work directory can contain log records, potentially including sensitive event data, held on disk before delivery. If the work directory is readable or writable by non-privileged accounts, buffered log data could be disclosed, altered, or deleted before it reaches its destination. Restricting the work directory to root protects buffered log data at rest.

Impact:

None. The recommended ownership and permissions match the rsyslog package default (0700 root:root), so applying this recommendation does not change default system behavior. If a non-default work directory is configured, the same ownership and mode should be applied to it.

Audit:

Verify that the rsyslog work directory is owned by root, group-owned by root, and mode 0700 or more restrictive. The work directory is set by the $WorkDirectory / global(workDirectory=...) directive and defaults to /var/lib/rsyslog.

1. Run the following command to determine the configured work directory:

```
# grep -Psi -- '^[^#]*(WorkDirectory|workDirectory)\s*[= ]' /etc/rsyslog.conf
/etc/rsyslog.d/*.conf 2>/dev/null
```

If a line is returned, use the path it sets in the next command; if nothing is returned, the default /var/lib/rsyslog applies.

2. Run the following command to list the work directory if it is not owned by root, not group-owned by root, or less restrictive than mode 0700 (substitute the configured path if $WorkDirectory is set to a non-default location):

```
# find /var/lib/rsyslog -maxdepth 0 -type d \( ! -user root -o ! -group root
-o -perm /0077 \) 2>/dev/null
```

Nothing should be returned.

Note: If rsyslog is not installed, the work directory may not exist; the command returns nothing and this recommendation is Not Applicable.

Remediation:

Run the following commands to set ownership to root:root and remove all group and other access from the work directory (substitute a non-default $WorkDirectory path if configured):

```
# chown root:root /var/lib/rsyslog
# chmod go-rwx /var/lib/rsyslog
```

Default Value:

/var/lib/rsyslog is mode 0700, owner root, group root (EL9 rsyslog package default).

References:

1. GRID: NIX-00000428
2. https://www.rsyslog.com/doc/rainerscript/global.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.	●	●	●

6.1.2.7 Ensure rsyslog is not configured as a log server (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

rsyslog can be configured to receive log messages from remote hosts over the network by loading a remote-input module (imtcp, imudp, imptcp, or imrelp) and starting a listener, which makes the host act as a central log server.

Rationale:

A general-purpose host that is not intended to be a central log collector should not run a syslog listener. Accepting remote log messages opens a network-facing service that increases the attack surface and exposes the host to remote log injection or poisoning (which can corrupt the local log store) and to resource exhaustion from unsolicited traffic. Only hosts explicitly designated and hardened as log servers should receive remote logs.

Impact:

None for a host that is not intended to receive remote logs. A host intentionally deployed as a central log server is outside the scope of this recommendation and should not apply it.

Audit:

Review the effective rsyslog configuration and verify the system is not configured to accept incoming logs from remote hosts. A log server may be configured with any of the input modules `imtcp`, `imudp`, `imptcp`, or `imrelp`. Run the following script (it assembles the effective configuration, including non-standard includes, then checks for remote-input modules in both the modern and legacy syntax):

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

{
   a_conf=(/etc/rsyslog.conf /etc/rsyslog.d/*.conf)
   while IFS= read -r l_inc; do
      for l_f in $l_inc; do
         [ -d "$l_f" ] && a_conf+=("$l_f"/*.conf) || a_conf+=("$l_f")
      done
   done < <(grep -Psoih --
'^\h*(?:\$IncludeConfig\h+|include\h*\(\h*file\h*=\h*"?)\K[^")\h]+'
"${a_conf[@]}" 2>/dev/null)
   grep -Psi -- '^\h*module\(load="?(imtcp|imudp|imptcp|imrelp)"?\)'
"${a_conf[@]}" 2>/dev/null
   grep -Psi -- '^\h*input\(type="?(imtcp|imudp|imptcp|imrelp)"?'
"${a_conf[@]}" 2>/dev/null
   grep -Psi -- '^\h*\$ModLoad\h+(imtcp|imudp|imptcp|imrelp)\b'
"${a_conf[@]}" 2>/dev/null
   grep -Psi --
'^\h*\$(InputTCPServerRun|UDPServerRun|InputPTCPServerRun|InputRELPServerRun)
\b' "${a_conf[@]}" 2>/dev/null
}
```

Nothing should be returned. If any output is returned, a remote-input module or network listener is configured and the system is acting as a log server, which does not meet this recommendation.

Remediation:

1. Should any active log server configuration be found by the audit, edit the file(s) where those directives appear - `/etc/rsyslog.conf`, a drop-in under `/etc/rsyslog.d/`, or any non-standard file identified by the audit's configuration discovery - and remove or comment out the remote-input module load, input, and run directives for all of `imtcp`, `imudp`, `imptcp`, and `imrelp`. Ensure none of the following entries are present:

advanced format

```
module(load="imtcp")
module(load="imudp")
module(load="imptcp")
module(load="imrelp")
input(type="imtcp" port="514")
input(type="imudp" port="514")
input(type="imptcp" port="514")
input(type="imrelp" port="2514")
```

deprecated legacy format

```
$ModLoad imtcp
$ModLoad imudp
$ModLoad imptcp
$ModLoad imrelp
$InputTCPServerRun 514
$UDPServerRun 514
$InputPTCPServerRun 514
$InputRELPServerRun 2514
```

2. Run the following command to restart `rsyslog`:

```
# 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.		●	●

6.1.2.8 Ensure rsyslog logging rules are configured (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The /etc/rsyslog.conf file and the drop-in files ending in .conf under /etc/rsyslog.d/ specify rsyslog's logging rules - which classes of messages (by facility and severity) are recorded and which files they are written to.

Rationale:

Security-relevant events - such as authentication successes and failures, privilege escalation, and root activity - are recorded through rsyslog. If logging rules are not configured to capture and retain these events, the records needed for accountability, detection, and investigation may never be written.

Audit:

1. Run the following script to assemble the effective rsyslog configuration and review the active logging rules; verify appropriate logging is set per site policy:

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

{
   a_conf=(/etc/rsyslog.conf /etc/rsyslog.d/*.conf)
   while IFS= read -r l_inc; do
      for l_f in $l_inc; do
         [ -d "$l_f" ] && a_conf+=("$l_f"/*.conf) || a_conf+=("$l_f")
      done
   done < <(grep -Psoih --
'^\h*(?:\$IncludeConfig\h+|include\h*\(\h*file\h*=\h*"?)\K[^")\h]+'
"${a_conf[@]}" 2>/dev/null)
   grep -Ph -- '^\h*[^#[:space:]]' "${a_conf[@]}" 2>/dev/null
}
```

The active logging rules across the effective configuration are listed. Rules should route the security-relevant facilities (for example authpriv, auth, cron, mail) to log files under /var/log/ per site policy.

Example output:

```
module(load="imuxsock")
global(workDirectory="/var/lib/rsyslog")
include(file="/etc/rsyslog.d/*.conf" mode="optional")
*.info;mail.none;authpriv.none;cron.none    /var/log/messages
authpriv.*                                  /var/log/secure
mail.*                                      -/var/log/maillog
cron.*                                      /var/log/cron
*.emerg                                     :omusrmsg:*
```

2. Run the following command to verify the expected log files exist and are being written to:

```
# ls -l /var/log/
```

Example output:

```
-rw-------.  1 root root  102400 Jul 31 10:15 messages
-rw-------.  1 root root   20480 Jul 31 10:15 secure
-rw-------.  1 root root    8192 Jul 31 10:14 cron
-rw-------.  1 root root    4096 Jul 31 10:10 maillog
```

Recent modification times on the security-relevant log files indicate rsyslog is actively writing to them.

Remediation:

1. 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/secure
mail.*                                  -/var/log/mail
mail.info                               -/var/log/mail.info
mail.warning                            -/var/log/mail.warn
mail.err                                /var/log/mail.err
cron.*                                  /var/log/cron
*.=warning;*.=err                       -/var/log/warn
*.crit                                  /var/log/warn
*.*;mail.none;news.none                 -/var/log/messages
local0,local1.*                         -/var/log/localmessages
local2,local3.*                         -/var/log/localmessages
local4,local5.*                         -/var/log/localmessages
local6,local7.*                         -/var/log/localmessages
```

2. Run the following command to restart `rsyslog`:

```
# 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.		●	●

6.1.2.9 Ensure rsyslog is configured to send logs to a remote log host (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

rsyslog can forward the log events it collects to a remote log host over the network (for example, with an omfwd action), sending a copy of local log data to a central collector for 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.

Impact:

Forwarding introduces a dependency on the remote collector and the network path to it; it does not replace local logging. Forwarded log volume consumes network bandwidth and collector storage, which should be sized accordingly.

Audit:

Run the following script to assemble the effective rsyslog configuration and verify logs are forwarded to a central log host. Both the basic and advanced configuration formats are accepted:

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

{
   a_conf=(/etc/rsyslog.conf /etc/rsyslog.d/*.conf)
   while IFS= read -r l_inc; do
      for l_f in $l_inc; do
         [ -d "$l_f" ] && a_conf+=("$l_f"/*.conf) || a_conf+=("$l_f")
      done
   done < <(grep -Psoih --
'^\h*(?:\$IncludeConfig\h+|include\h*\(\h*file\h*=\h*"?)\K[^")\h]+'
"${a_conf[@]}" 2>/dev/null)
   grep -Psi -- '^\h*[^#]\S*\h+@@\S+' "${a_conf[@]}" 2>/dev/null
   grep -Psi -- '^\s*([^#]+\s+)?action\(([^#]+\s+)?\btarget="?[^#"]+"?\b'
"${a_conf[@]}" 2>/dev/null
}
```

A forwarding line naming the remote loghost should be returned - either a basic-format line or an advanced `omfwd` action.

Example output:

```
*.* @@loghost.example.com
*.* action(type="omfwd" target="loghost.example.com" port="514"
protocol="tcp")
```

Note: `@@` denotes a TCP connection; a single `@` denotes UDP, which does not meet this recommendation.

Remediation:

1. Edit the `/etc/rsyslog.conf` and `/etc/rsyslog.d/*.conf` files and add a forwarding action (where `loghost.example.com` is the name of your central log host; the `target` directive may be a fully qualified domain name or an IP address):

```
*.* action(type="omfwd" target="loghost.example.com" port="514"
protocol="tcp"
           action.resumeRetryCount="100"
           queue.type="LinkedList" queue.size="1000")
```

2. Run the following command to restart `rsyslog`:

```
# systemctl restart rsyslog
```

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.		●	●

6.1.2.10 Ensure a rsyslog TLS network stream driver is installed (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

rsyslog encrypts forwarded syslog traffic using a TLS network stream driver, a layer between rsyslog and the transport that provides confidentiality, integrity, and peer authentication for messages sent over the network using the syslog-transport-tls standard. Two interchangeable drivers are available and packaged separately: the GnuTLS driver gtls (package rsyslog-gnutls) and the OpenSSL driver ossl (package rsyslog-openssl). Either driver satisfies this recommendation; only the selected driver's package needs to be installed.

Installing a TLS network stream driver ensures the capability to forward logs over an encrypted, authenticated channel is available whenever remote forwarding is used.

Rationale:

Traditional syslog forwarding is a clear-text protocol, so anyone able to observe the network can read forwarded log data. A TLS network stream driver implements the SSL/TLS protocols that encrypt syslog communication in transit and authenticate the remote peer. The driver package must be installed before rsyslog can select a TLS driver, so installing rsyslog-gnutls or rsyslog-openssl is the prerequisite for any encrypted forwarding configuration.

Impact:

Installing a TLS driver package enables, but does not by itself activate, encrypted forwarding; it must be paired with driver, CA, and (for mutual authentication) certificate configuration. Once TLS forwarding is enabled, ongoing certificate lifecycle management is required: an expired or misconfigured certificate or trust path will cause forwarding to fail until the trust chain is restored.

Audit:

Verify that a rsyslog TLS network stream driver package is installed. Either the GnuTLS driver (rsyslog-gnutls) or the OpenSSL driver (rsyslog-openssl) satisfies this recommendation.

Run the following command to verify that either the rsyslog-gnutls or the rsyslog-openssl package is installed:

```
# rpm -q rsyslog-gnutls rsyslog-openssl
```

Verify at least one of the two driver packages is installed.

Example output:

```
rsyslog-gnutls-<version>
package rsyslog-openssl is not installed
```

It is normal for the unused driver to report that it is not installed.

Remediation:

Install a rsyslog TLS network stream driver. Only one driver package is required; the GnuTLS driver is the recommended default.

Run the following command to install the `rsyslog-gnutls` package:

```
# dnf install rsyslog-gnutls
```

Or, if the OpenSSL driver is preferred, run the following command to install the `rsyslog-openssl` package:

```
# dnf install rsyslog-openssl
```

Default Value:

Not installed. rsyslog is installed by default; the rsyslog-gnutls and rsyslog-openssl TLS driver subpackages are not.

References:

1. GRID: NIX-00000281
2. https://www.rsyslog.com/doc/concepts/netstrm_drvr.html
3. https://www.rsyslog.com/doc/tutorials/tls.html

6.1.2.11 Ensure rsyslog network forwarding uses TLS (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

When rsyslog forwards logs to a remote host, the omfwd output action can wrap the connection in TLS by selecting a TLS network stream driver (gtls or ossl), running that driver in TLS-only mode, and setting a peer authentication mode. The driver may be selected globally (DefaultNetstreamDriver) or per action (StreamDriver), using either the modern advanced (global(...) / action(...)) syntax or the legacy $-directive syntax. This recommendation applies only where rsyslog forwards to a remote host; hosts that do not forward syslog over the network may document it as Not Applicable.

Rationale:

If forwarding is not wrapped in TLS, syslog traffic crosses the network in clear text and is subject to interception, modification, or silent loss. Selecting a TLS stream driver and forcing TLS-only mode ensures the connection is encrypted, and setting an authentication mode ensures rsyslog will only deliver logs to a verified remote peer, defeating passive sniffing and man-in-the-middle attacks against the log stream.

Impact:

Enabling TLS forwarding requires that the remote collector also accept TLS on the encrypted port and that the necessary CA (and, for mutual authentication, client certificate and key) material be present and current. Using TLS-only mode means forwarding will not silently fall back to clear text: if the TLS handshake fails, messages queue or drop rather than being sent unencrypted, so certificate or connectivity problems become forwarding outages until corrected.

Audit:

Verify that rsyslog forwards logs over TLS. A compliant forwarding action selects a TLS network stream driver (`gtls` or `ossl`), runs it in TLS-only mode, sets a peer authentication mode, and targets the encrypted syslog-over-TLS port (`6514`). All four elements must be present on the same forwarding action. The checks accept either driver and both the modern advanced syntax (`global(...)` / `action(...)`) and the legacy `$`-directive syntax; legacy directives are deprecated and should be migrated to the advanced format.

1. Identify all active rsyslog configuration files. In addition to `/etc/rsyslog.conf` and the `.conf` drop-ins under `/etc/rsyslog.d/`, rsyslog may include further files via `$IncludeConfig` or `include()`. Run the following command and include any files it returns in the searches below (the bound automated check assembles this set recursively):

```
# grep -Psoih --
'^\h*(?:\$IncludeConfig\h+|include\h*\(\h*file\h*=\h*"?)\K[^")\h]+'
/etc/rsyslog.conf /etc/rsyslog.d/*.conf
```

On a system that uses only the standard configuration paths, nothing is returned.

2. Run the following command to verify a TLS network stream driver (`gtls` or `ossl`) is selected:

```
# grep -Psi -- '^[^#]*(StreamDriver|DefaultNetstreamDriver)(\.Name)?\s*[=
]\s*"?(gtls|ossl)"?' /etc/rsyslog.conf /etc/rsyslog.d/*.conf
```

A line setting the driver to `gtls` or `ossl` should be returned (for example `StreamDriver="gtls"`, or the legacy `$DefaultNetstreamDriver gtls`).

3. Run the following command to verify the driver runs in TLS-only mode (mode `1`):

```
# grep -Psi --
'^[^#]*(StreamDriverMode\s*=\s*"?1"?|ActionSendStreamDriverMode\s+1)'
/etc/rsyslog.conf /etc/rsyslog.d/*.conf
```

A line setting the driver mode to `1` should be returned (for example `StreamDriverMode="1"`, or the legacy `$ActionSendStreamDriverMode 1`).

4. Run the following command to verify a peer authentication mode is set:

```
# grep -Psi --
'^[^#]*(StreamDriverAuthMode\s*=\s*"?(x509/name|x509/certvalid|x509/fingerpri
nt|anon)"?|ActionSendStreamDriverAuthMode\s+(x509/name|x509/certvalid|x509/fi
ngerprint|anon))' /etc/rsyslog.conf /etc/rsyslog.d/*.conf
```

A line setting the authentication mode should be returned (for example `StreamDriverAuthMode="x509/name"`, or the legacy `$ActionSendStreamDriverAuthMode x509/name`). `x509/name` is recommended because it validates the remote certificate chain and matches the peer identity; `anon` encrypts but does not authenticate the peer and should be avoided in production.

5. Run the following command to verify the forwarding action targets the encrypted syslog-over-TLS port (`6514`):

```
# grep -Psi -- '^[^#]*(port\s*=\s*"?6514"?|:6514)' /etc/rsyslog.conf
/etc/rsyslog.d/*.conf
```

A line targeting port `6514` should be returned (for example `port="6514"`, or the legacy `@@loghost:6514`).

All four elements (steps 2-5) must be present on the same forwarding action.

Remediation:

Configure rsyslog to forward over TLS using a drop-in file, the recommended method for local rsyslog configuration, and ensure no plaintext forwarding action remains.

1. Run the following command to locate any existing forwarding actions for review:

```
# grep -PHns -- 'type="omfwd"|^\h*[^#]*@' /etc/rsyslog.conf
/etc/rsyslog.d/*.conf
```

2. In each file returned, remove or convert to TLS any forwarding action that is not already TLS-secured, so the only remaining forwarding path uses TLS.
3. Create a drop-in file such as /etc/rsyslog.d/60-forward-tls.conf with the TLS forwarding configuration. The example below uses the GnuTLS driver:

```
# Select the TLS driver and trust material globally
global(
  DefaultNetstreamDriver="gtls"
  DefaultNetstreamDriverCAFile="/etc/pki/tls/certs/ca-bundle.crt"
)

# Forward all messages to the remote collector over TLS only
*.* action(type="omfwd"
           target="loghost.example.com" port="6514" protocol="tcp"
           StreamDriver="gtls"
           StreamDriverMode="1"
           StreamDriverAuthMode="x509/name"
           StreamDriverPermittedPeers="loghost.example.com"
           action.resumeRetryCount="100"
           queue.type="LinkedList" queue.size="1000")
```

4. Run the following command to restart `rsyslog`:

```
# systemctl restart rsyslog
```

If the OpenSSL driver is preferred, substitute ossl for gtls in the DefaultNetstreamDriver and StreamDriver settings above.

Default Value:

No TLS driver configured; forwarding, if any, is plain TCP or UDP in clear text.

References:

1. GRID: NIX-00000274
2. https://www.rsyslog.com/doc/tutorials/tls_cert_client.html
3. https://www.rsyslog.com/doc/concepts/netstrm_drvr.html
4. https://www.rsyslog.com/doc/configuration/modules/omfwd.html

6.1.2.12 Ensure rsyslog TLS CA certificate is configured (Manual)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

When rsyslog forwards logs over TLS, it uses a Certificate Authority (CA) certificate to verify the remote collector's TLS certificate. The CA is specified with the global DefaultNetstreamDriverCAFile directive, or per action with StreamDriver.CAFile (rsyslog 8.2108.0 and later). For mutual authentication, the client also presents its own certificate and key (DefaultNetstreamDriverCertFile / DefaultNetstreamDriverKeyFile, or the per-action StreamDriver.CertFile / StreamDriver.KeyFile). This recommendation applies where rsyslog forwards over TLS.

Rationale:

A trusted CA certificate ensures rsyslog only connects to log servers whose TLS certificates are signed by a known, trusted authority. Without a properly configured CA, forwarding can proceed without verifying the server's identity, exposing the host to man-in-the-middle attacks that intercept, tamper with, or silently discard forwarded logs, resulting in undetected data leakage or loss of audit-trail integrity.

Impact:

A CA certificate (and, for mutual authentication, a client certificate and key) must be provisioned and kept current. If the referenced file is missing, expired, or points to the wrong path, TLS forwarding fails until the trust chain is restored. On systems with SELinux enforcing, a CA file placed outside the paths rsyslogd is allowed to read will be blocked even when ownership and permissions are correct, so custom certificate locations require an SELinux file-context adjustment in addition to the rsyslog configuration.

Audit:

Verify that rsyslog is configured with a trusted CA certificate for TLS forwarding. The CA certificate may be set globally with `DefaultNetstreamDriverCAFile` or per forwarding action with `StreamDriver.CAFile` (rsyslog 8.2108.0 and later); either form satisfies this recommendation. Both the modern advanced syntax and the legacy `$`-directive syntax are accepted; legacy directives are deprecated and should be migrated when possible.

1. Identify all active rsyslog configuration files. In addition to `/etc/rsyslog.conf` and the `.conf` drop-ins under `/etc/rsyslog.d/`, rsyslog may include further files via `$IncludeConfig` or `include()`. Run the following command and include any files it returns in the search below:

```
# grep -Psoih --
'^\h*(?:\$IncludeConfig\h+|include\h*\(\h*file\h*=\h*"?)\K[^")\h]+'
/etc/rsyslog.conf /etc/rsyslog.d/*.conf
```

On a system that uses only the standard configuration paths, nothing is returned.

2. Run the following command to verify a CA certificate file is configured:

```
# grep -Psi --
'^[^#]*(DefaultNetstreamDriverCAFile|StreamDriver\.CAFile)\s*[= ]'
/etc/rsyslog.conf /etc/rsyslog.d/*.conf
```

A line naming a CA certificate file should be returned (for example `DefaultNetstreamDriverCAFile="/etc/pki/tls/certs/ca-bundle.crt"`).

3. For each CA file configured in step 2, run the following command to confirm the file exists and is a readable certificate (substitute the path returned above):

```
# openssl x509 -in /etc/pki/tls/certs/ca-bundle.crt -noout -subject
```

A certificate subject line should be returned (for example `subject=C = US, O = Example, CN = Example Logging CA`). If the file is missing or unreadable, or `openssl` reports an error, the CA certificate is not correctly configured.

Confirm the CA originates from a known, trusted authority appropriate to the environment and local site policy - a public CA or an organization-managed internal enterprise CA. Avoid ad-hoc or untrusted self-signed certificates in production; the steps above confirm a CA file is configured and readable but do not establish that the issuing authority is trusted.

Remediation:

Configure rsyslog with the CA certificate that signs the remote log server's TLS certificate.

1. Add the global CA directive to the drop-in file that defines the TLS forwarding action (for example `/etc/rsyslog.d/60-forward-tls.conf`):

```
global(DefaultNetstreamDriverCAFile="/etc/pki/tls/certs/ca-bundle.crt")
```

Set the path to the CA that signed the log server certificate: use the system trust bundle `/etc/pki/tls/certs/ca-bundle.crt` when the server certificate chains to a CA already in the system trust store, or a dedicated CA file (for example `/etc/pki/rsyslog/ca.pem`) for an internal certificate authority.

2. Run the following command to restart `rsyslog`:

```
# systemctl restart rsyslog
```

Note: On rsyslog 8.2108.0 and later the CA may instead be set per forwarding action with `StreamDriver.CAFile="..."` inside the `action(type="omfwd" ...)` block, preferred when different targets use different CA chains.

If a non-default CA path is used on a system with SELinux in enforcing mode, label the path so rsyslogd can read it, for example:

```
# semanage fcontext -a -t cert_t '/etc/pki/rsyslog(/.*)?'
# restorecon -Rv /etc/pki/rsyslog
```

Default Value:

No CA certificate configured.

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

Additional Information:

On EL-family systems the default MAC framework is SELinux, and rsyslogd runs confined in the rsyslogd_t domain. Certificates under the system trust store (/etc/pki/tls/certs/, /etc/pki/ca-trust/) are readable by default. A CA or key file placed elsewhere may be denied by SELinux even with correct ownership and permissions, producing TLS handshake failures rather than an obvious permission error. Diagnose with ausearch -m avc -c rsyslogd or journalctl (look for AVC ... denied ... comm="rsyslogd"), and relabel custom paths with semanage fcontext + restorecon as shown in the remediation.

6.1.2.13 Ensure rsyslog TLS client certificate is configured (Manual)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

In mutual (two-way) TLS, the rsyslog forwarding client presents its own X.509 certificate and private key so the receiving log server can authenticate it, in addition to the client verifying the server against a trusted CA. The client certificate and key are configured with the modern advanced parameters `DefaultNetstreamDriverCertFile` / `DefaultNetstreamDriverKeyFile` (or the per-action `StreamDriver.CertFile` / `StreamDriver.KeyFile`), or the equivalent legacy `$DefaultNetstreamDriverCertFile` / `$DefaultNetstreamDriverKeyFile` directives.

Rationale:

With one-way TLS the forwarding client verifies the receiver, but the receiver cannot verify the forwarder: any host that can reach the collector and present a certificate chaining to a trusted CA - or connect with an anonymous authentication mode - can establish a session and inject log records. Configuring a client certificate enables mutual authentication so the receiving log server can confirm the identity of each connecting forwarder, preventing unauthorized or spoofed hosts from injecting or poisoning the central log store.

Impact:

Mutual TLS requires a client certificate and key on each forwarding host and the corresponding client identity to be authorized on the receiving log server, and those certificates must be maintained and rotated. A missing, mismatched, or expired client certificate will prevent the forwarder from connecting, interrupting log delivery until it is corrected.

Audit:

This recommendation applies only where rsyslog forwards logs over TLS.

Verify that rsyslog presents a client certificate and private key for mutual TLS authentication. The certificate and key may be set globally (`DefaultNetstreamDriverCertFile` / `DefaultNetstreamDriverKeyFile`) or per forwarding action (`StreamDriver.CertFile` / `StreamDriver.KeyFile`); either form satisfies this recommendation. Legacy `$`-directive forms are also accepted but are deprecated and should be migrated to the advanced format.

1. Identify all active rsyslog configuration files. In addition to `/etc/rsyslog.conf` and the `.conf` drop-ins under `/etc/rsyslog.d/`, rsyslog may include further files via `$IncludeConfig` or `include()`. Run the following command and include any files it returns in the searches below:

```
# grep -Psoih --
'^\h*(?:\$IncludeConfig\h+|include\h*\(\h*file\h*=\h*"?)\K[^")\h]+'
/etc/rsyslog.conf /etc/rsyslog.d/*.conf
```

On a system that uses only the standard configuration paths, nothing is returned.

2. Run the following command to verify a client certificate file is configured:

```
# grep -Psi -- '(DefaultNetstreamDriverCertFile|StreamDriver\.CertFile)\h*[=
]' /etc/rsyslog.conf /etc/rsyslog.d/*.conf
```

A line naming a client certificate file should be returned (for example `DefaultNetstreamDriverCertFile="/etc/pki/tls/certs/rsyslog-client.crt"`).

3. Run the following command to verify a client private key file is configured:

```
# grep -Psi -- '(DefaultNetstreamDriverKeyFile|StreamDriver\.KeyFile)\h*[= ]'
/etc/rsyslog.conf /etc/rsyslog.d/*.conf
```

A line naming a private key file should be returned (for example `DefaultNetstreamDriverKeyFile="/etc/pki/tls/private/rsyslog-client.key"`).

4. For the certificate and key files configured in steps 2 and 3, run the following command to confirm each file exists and to check the private key's ownership and access (substitute the paths returned above):

```
# ls -l /etc/pki/tls/certs/rsyslog-client.crt /etc/pki/tls/private/rsyslog-
client.key
```

Both files should exist. The private key file should be owned by `root` and not readable by group or other (for example `-rw------- 1 root root`).

Confirm the client certificate is issued by a certificate authority the receiving log server trusts and matches the forwarding host's identity.

Remediation:

Provision a client certificate and private key issued by a CA the receiving log server trusts, then configure rsyslog to present them.

1. Add the following to a drop-in such as `/etc/rsyslog.d/60-forward-tls.conf`:

```
global(
  DefaultNetstreamDriver="gtls"
  DefaultNetstreamDriverCAFile="/etc/pki/tls/certs/ca-bundle.crt"
  DefaultNetstreamDriverCertFile="/etc/pki/tls/certs/rsyslog-client.pem"
  DefaultNetstreamDriverKeyFile="/etc/pki/tls/private/rsyslog-client.key"
)
```

2. Run the following command to restrict permissions on the private key:

```
# chmod 0600 /etc/pki/tls/private/rsyslog-client.key
```

3. Run the following command to restart `rsyslog`:

```
# systemctl restart rsyslog
```

On the receiving log server, require and authorize client certificates (for example `StreamDriverAuthMode="x509/name"` with an appropriate `permittedPeer` list) so only authenticated forwarders may connect.

References:

1. GRID: NIX-00000429

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.		●	●

6.1.2.14 Ensure logrotate is configured (Manual)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

logrotate manages the automatic rotation, compression, and removal of log files so they do not grow without bound. The log files produced by rsyslog are rotated according to /etc/logrotate.conf and the drop-in files under /etc/logrotate.d/ (for example /etc/logrotate.d/syslog).

Rationale:

Without log rotation, log files grow without bound and can consume all available space on their filesystem, disrupting logging itself and other services that share that filesystem. Rotating and aging logs bounds disk use while retaining recent records for the period required by site policy.

Impact:

Enforcing log rotation ages out and removes log files once they pass the configured retention window. If the rotation and retention settings are more aggressive than the organization's requirements, historical logs needed for investigation or compliance may be removed prematurely; retention should be set to satisfy the applicable log-retention policy.

Audit:

Run the following command to review the configured log rotation settings across the main configuration file and the drop-in files:

```
# grep -PH '^\h*(daily|weekly|monthly|rotate|maxage|compress|create)\b'
/etc/logrotate.conf /etc/logrotate.d/*
```

The output should show a rotation schedule (for example `daily`, `weekly`, or `monthly`), a bounded retention count (`rotate`), and, where required by site policy, a maximum age (`maxage`) and compression (`compress`). Verify the settings meet the organization's log-retention policy.

Example output:

```
/etc/logrotate.conf:weekly
/etc/logrotate.conf:rotate 13
/etc/logrotate.conf:maxage 90
/etc/logrotate.conf:compress
/etc/logrotate.conf:create
```

Remediation:

Edit `/etc/logrotate.conf` (and any files under `/etc/logrotate.d/`) to rotate logs on a fixed schedule, retain a bounded number of rotations, compress old logs, and set a maximum age. Apply the following baseline in `/etc/logrotate.conf`:

```
# rotate log files weekly
weekly

# keep 13 weeks of backlogs
rotate 13

# remove rotated logs older than 90 days
maxage 90

# compress rotated logs
compress

# create new (empty) log files after rotating old ones
create
```

Note: Tune the `weekly`, `rotate`, and `maxage` values to your organization's log-retention policy; the values above are a recommended starting point.

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.

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.		●	●

6.1.3 Configure Logfiles

This section provides recommendations for restricting access to log files so their contents cannot be read, modified, or removed by unauthorized users. Log files are typically stored under `/var/log/`, though applications may write elsewhere; they record information from the many services on the local system - and, where a centralized log server is used, from other systems - and frequently contain sensitive information, so their access must be controlled. Access requirements vary by log class: sensitive system and authentication logs are held most strictly, while login-accounting, journal, and package or cloud logs follow their platform defaults.

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 local system, or in the event of a centralized log server, other systems' logs as well.

In general log files are found in `/var/log/`, although applications can be configured to store logs elsewhere. Should your application store logs in another location, apply the same audit to that location.

Rationale:

It is important that log files have the correct permissions to ensure that sensitive data is protected and that only the appropriate users / groups have access to them.

Impact:

Tightening permissions and ownership on files under `/var/log/` can disrupt services or tooling that rely on broader access. The remediation is applied per file type and is non-destructive, but capture its output in case a service's logging needs to be reconfigured.

Audit:

Log files under `/var/log/` do not all share the same expected access: login-accounting, systemd journal, and package or cloud logs each have different owner, group, and mode requirements from the general system logs. The intent is to keep sensitive log data - the system and authentication logs - from being read or altered by unauthorized users, while accepting each platform's own default access for non-sensitive logs so that compliance stays stable across log rotation and reboot.

A file is compliant when its mode is the value shown (or more restrictive) and its owner and group are among those listed for its class:

Log file(s)	Owner	Group	Mode (or more restrictive)
General system logs - `messages`, `syslog`, `secure*`, `auth.log`, and any other `/var/log` log not covered below	`root` or `syslog`	`root` or `adm`	`0640`
Login accounting - `wtmp*`, `btmp*`, `lastlog*`	`root`	`root` or `utmp`	`0664`
systemd journal - `*.journal`, `*.journal~`	`root`	`root` or `systemd-journal`	`0640`
Package-manager logs - files written by the system package manager: `/var/log/apt/*` (apt), `/var/log/dnf.*`, `/var/log/dnf5.log`, `/var/log/hawkey.log*` (dnf/dnf5), `/var/log/zypp/*` and `/var/log/zypper.log*` (zypper), `/var/log/yum.log*` (legacy yum)	`root`	`root` or `adm`	`0644`
Cloud or agent logs - `cloud-init.log*`, `localmessages*`, `waagent.log*`	`root` or `syslog`	`root` or `adm`	`0644`
Service-account logs - a log written by a dedicated service account (an account whose login shell is not listed in `/etc/shells`)	that account, `root`, or `syslog`	that account's group, `root`, or `adm`	`0640`

Run the following command to list every file under `/var/log/` whose mode, owner, or group differs from the strictest baseline (mode `0640`, owner `root`, group `root`); these are the files to compare against the table. The output columns are the file name, owner, group, and octal mode - matching the table order:

```
# find -L /var/log -type f \( -perm /0137 -o ! -user root -o ! -group root \)
-exec stat -Lc '%n %U %G %#a' {} +
```

For each file returned, find its class in the table and confirm its mode is the listed value (or more restrictive) and its owner and group are among those listed. Files that are not returned are already mode `0640` or more restrictive, owned by `root`, and group-owned by `root`, and need no further review.

Example output (`wtmp` is returned because it is group `utmp` and mode `0664` - the expected access for the login-accounting class, so it is compliant; a general system log returned at `0644` would not be compliant, because it would be world-readable):

```
/var/log/wtmp root utmp 0664
```

Any returned file whose mode, owner, or group does not match its class does not meet this recommendation. Application logs stored outside `/var/log/` should be held to the equivalent standard for their class.

Remediation:

Log file access is most reliably enforced at the point of creation: if the mechanism that creates each log sets the correct mode and ownership, new and rotated logs are born compliant and the state does not drift back on the next rotation or reboot. Correct the creating mechanism first, then adjust any existing files.

1. Configure the log-creating mechanisms so new logs are created with restrictive access:
    - For the general system and authentication logs written by rsyslog, set the file creation mode to `0640` or more restrictive (see the recommendation for configuring the rsyslog log file creation mode).
    - systemd journal files are created mode `0640`, group `systemd-journal` by systemd-journald; no change is required.
    - For logs rotated by logrotate, use a `create` directive with the correct mode and ownership so rotated files retain the restrictive access (see the recommendation for configuring logrotate).
2. Correct existing files. For each file returned by the Audit, set its mode and ownership to match its class in the table. For a general system or authentication log that is more permissive than `0640`, or not owned by `root:root`, run the following commands to restrict its access and set its ownership:

```
# chmod u-x,g-wx,o-rwx <file>
# chown root:root <file>
```

For the other classes, apply the corresponding values instead:

- Login accounting (`wtmp*`, `btmp*`, `lastlog*`) - mode `0664` or more restrictive, group `utmp`.
- systemd journal (`*.journal`, `*.journal~`) - group `systemd-journal`.
- Package-manager logs (files written by the system package manager, e.g. `/var/log/apt/`, `/var/log/dnf.*`, `/var/log/dnf5.log`, `/var/log/zypp/`) - mode `0644` or more restrictive, group `root` or `adm`.
- Cloud or agent logs (`cloud-init.log*`, `localmessages*`, `waagent.log*`) - mode `0644` or more restrictive.
- Service-account logs - keep that account as the owner and group; only remove world and any excess access. Do not force `root` ownership, as it can break the service that writes the log.

3. Re-run the Audit to confirm no files remain.

Note: Each class's ceiling matches the platform's default creation mode, so a file corrected to its class value stays compliant when the service or logrotate recreates it - provided the creating mechanism in step 1 is set correctly. Applying step 2 without step 1 may allow a sensitive log to be recreated world-readable on the next rotation.

References:

1. GRID: NIX-00000262
2. NIST SP 800-53 :: AC-3, AU-9

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:

- On the command line using the `auditctl` utility. These rules are not persistent across reboots.
- In `/etc/audit/audit.rules`. These rules have to be merged and loaded before they are active.

Notes:

- For 64 bit systems that have `arch` as a rule parameter, you will need two rules: one for 64 bit and one for 32 bit systems calls. For 32 bit systems, only one rule is needed.
- If the auditing system is configured to be locked (`-e 2`), a system reboot will be required in order to load any changes.
- Key names are optional on the rules and will not be used in compliance auditing. The usage of key names is highly recommended as it facilitates organization and searching; as such, all remediation steps will have key names supplied.
- It is best practice to store the rules, in number prepended files, in `/etc/audit/rules.d/`. Rules must end in a `.rules` suffix. This then requires the use of `augenrules` to merge all the rules into `/etc/audit/audit.rules` based on their alphabetical (lexical) sort order. All benchmark recommendations follow this best practice for remediation, specifically using the prefix of `50` which is center weighed if all rule sets make use of the number prepending naming convention.
- Your system may have been customized to change the default `UID_MIN`. All sample output uses `1000`, but this value will not be used in compliance auditing. To confirm the `UID_MIN` for your system, run the following command: `awk '/^\s*UID_MIN/{print $2}' /etc/login.defs`

Normalization The Audit system normalizes some entries, so when you look at the sample output keep in mind that:

- With regards to users whose login UID is not set, the values `-1` / `unset` / `4294967295` are equivalent and normalized to `-1`.
- When comparing field types and both sides of the comparison is valid fields types, such as `euid!=uid`, then the auditing system may normalize such that the output is `uid!=euid`.
- Some parts of the rule may be rearranged whilst others are dependent on previous syntax. For example, the following two statements are the same:

```
-a always,exit -F arch=b64 -S execve -C uid!=euid -F auid!=-1 -F
key=user_emulation
```

and

```
-a always,exit -F arch=b64 -C euid!=uid -F auid!=unset -S execve -k
user_emulation
```

Capacity planning

The recommendations in this section implement auditing policies that not only produce large quantities of logged data, but may also negatively impact system performance. Capacity planning is critical in order not to adversely impact production environments.

- Disk space. If a significantly large set of events are captured, additional on system or off system storage may need to be allocated. If the logs are not sent to a remote log server, ensure that log rotation is implemented else the disk will fill up and the system will halt. Even when logs are sent to a log server, ensure sufficient disk space to allow caching of logs in the case of temporary network outages.
- Disk IO. It is not just the amount of data collected that should be considered, but the rate at which logs are generated.
- CPU overhead. System call rules might incur considerable CPU overhead. Test the systems open/close syscalls per second with and without the rules to gauge the impact of the rules.

6.2.1 Configure auditd Service

This subsection covers baseline configuration of the Linux Audit daemon (`auditd`) service: installing the `audit` and `audit-libs` packages, enabling auditing early in the boot process (`audit=1`), setting an adequate kernel audit backlog limit (`audit_backlog_limit=`), and ensuring the `auditd` service is enabled and active.

The `audit=1` and `audit_backlog_limit=` settings are kernel boot parameters. These recommendations assume the grub2 bootloader and use `grubby` to read the effective boot entries, with `/etc/default/grub` as the persistent source; if a different bootloader is in use, apply the equivalent settings for that loader. Changes to kernel boot parameters take effect on the next reboot.

6.2.1.1 Ensure auditd packages are installed (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

`audit` is the userspace component of the Linux Auditing System. It is responsible for writing audit records to 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.

Impact:

With `auditd` installed and running, the system records audit events, which consumes some CPU and disk space for the audit logs; size audit log storage and retention accordingly.

Audit:

Run the following command to verify that the `audit` and `audit-libs` packages are installed:

```
# rpm -q audit audit-libs

audit-<version>
audit-libs-<version>
```

Remediation:

Run the following command to install `audit` and `audit-libs`:

```
# dnf install audit audit-libs
```

References:

1. GRID: NIX-00000300
2. NIST SP 800-53 :: AU-2, AU-12
3. NIST SP 800-53A :: AU-12.1 (ii)
4. RHEL 8 STIG Vul ID: V-230411
5. RHEL 8 STIG Rule ID: SV-230411r744000

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.	●	●	●

6.2.1.2 Ensure auditing for processes that start prior to auditd is enabled (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

The `audit=1` kernel boot parameter enables the Linux Auditing System early in the boot process, before the `auditd` service starts.

Rationale:

Audit events need to be captured on processes that start up prior to `auditd`, so that potential malicious activity cannot go undetected.

Impact:

Auditing processes that start before `auditd` increases the volume of audit records captured during early boot.

Audit:

Both of the following must be true.

1. Run the following command to verify the `audit=1` parameter is set in the running boot configuration:

```
# grubby --info=ALL | grep -Po '\baudit=1\b'

audit=1
```

`audit=1` may be returned multiple times.

2. Run the following command to verify the `audit=1` parameter is set in `/etc/default/grub`:

```
# grep -Psoi -- '^\h*GRUB_CMDLINE_LINUX="([^#\n\r]+\h+)?audit=1\b'
/etc/default/grub
```

Example output:

```
GRUB_CMDLINE_LINUX="quiet audit=1"
```

Remediation:

Run the following command to update the `grub2` configuration with `audit=1`:

```
# grubby --update-kernel ALL --args 'audit=1'
```

Edit `/etc/default/grub` and add `audit=1` to the `GRUB_CMDLINE_LINUX=` line between the opening and closing double quotes:

Example:

```
GRUB_CMDLINE_LINUX="quiet audit=1"
```

Note: `audit=1` is a kernel boot parameter; the change takes effect on the next reboot.

References:

1. GRID: NIX-00000302
2. NIST SP 800-53 :: AU-2, AU-12
3. NIST SP 800-53A :: AU-12.1 (ii)
4. RHEL 8 STIG GROUP ID: V-230468
5. RHEL 8 STIG Rule ID: SV-230468r792904

Additional Information:

Refer to the "Configure auditd Service" section overview for bootloader guidance (`grub2`, `grubby`, `/etc/default/grub`) and kernel boot-parameter reboot behavior.

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.	●	●	●

6.2.1.3 Ensure audit_backlog_limit is configured (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

The `audit_backlog_limit` parameter determines how many audit records can be held in the audit backlog.

Rationale:

During boot if `audit=1`, then the backlog will hold 64 records. If more than 64 records are created during boot, audit records will be lost and potential malicious activity could go undetected.

Impact:

A larger `audit_backlog_limit` reserves more kernel memory for the boot-time audit backlog buffer; the default of 64 may be too small and drop early-boot audit records.

Audit:

Both of the following must be true.

1. Run the following command to verify the `audit_backlog_limit=` parameter is set to an appropriate size for your organization:

```
# grubby --info=ALL | grep -Po "\baudit_backlog_limit=\d+\b"

audit_backlog_limit=<BACKLOG SIZE>
```

It is recommended that this value be `8192` or larger.

2. Run the following command to verify the `audit_backlog_limit=` parameter is set in `/etc/default/grub`:

```
# grep -Psoi --
'^\h*GRUB_CMDLINE_LINUX="([^#\n\r]+\h+)?\baudit_backlog_limit=\d+\b'
/etc/default/grub
```

Example output:

```
GRUB_CMDLINE_LINUX="quiet audit_backlog_limit=8192"
```

Remediation:

Run the following command to add `audit_backlog_limit=<BACKLOG SIZE>` to GRUB_CMDLINE_LINUX:

```
# grubby --update-kernel ALL --args 'audit_backlog_limit=<BACKLOG SIZE>'
```

Example:

```
# grubby --update-kernel ALL --args 'audit_backlog_limit=8192'
```

Edit `/etc/default/grub` and add `audit_backlog_limit=<BACKLOG SIZE>` to the `GRUB_CMDLINE_LINUX=` line between the opening and closing double quotes:

Example:

```
GRUB_CMDLINE_LINUX="quiet audit_backlog_limit=8192"
```

Note: `audit_backlog_limit` is a kernel boot parameter; the change takes effect on the next reboot.

Default Value:

`audit_backlog_limit=64`

References:

1. GRID: NIX-00000299
2. NIST SP 800-53 :: AU-2, AU-12
3. NIST SP 800-53A :: AU-12.1 (ii)

Additional Information:

Refer to the "Configure auditd Service" section overview for bootloader guidance (`grub2`, `grubby`, `/etc/default/grub`) and kernel boot-parameter reboot behavior.

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.	●	●	●

6.2.1.4 Ensure auditd service is enabled and active (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

The `auditd` service is the userspace daemon of the Linux Auditing System. When enabled and active it records system events to the audit log.

Rationale:

The capturing of system events provides system administrators with information to allow them to determine if unauthorized access to their system is occurring.

Impact:

Running `auditd` records system events, which consumes some CPU and disk space for the audit logs.

Audit:

Both of the following must be true.

1. Run the following command to verify the `auditd` service is enabled:

```
# systemctl is-enabled auditd

enabled
```

2. Run the following command to verify the `auditd` service is active:

```
# systemctl is-active auditd

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 :: AU-2, AU-12
3. NIST SP 800-53A :: AU-12.1 (ii)
4. RHEL 8 STIG Vul ID: V-244542
5. RHEL 8 STIG Rule ID: SV-244542r818838

Additional Information:

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

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.2 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.	●	●	●

6.2.2 Configure Data Retention

This subsection configures how the audit daemon (`auditd`) retains audit records on disk: the size of each audit log file (`max_log_file`), whether full logs are rotated and kept rather than deleted (`max_log_file_action`), and the actions taken when the partition holding the audit logs is full or low on space (`disk_full_action`, `disk_error_action`, `space_left_action`, `admin_space_left_action`).

By default `auditd` limits each audit log file to 8 MB (`max_log_file`) and rotates through a set number of files (`num_logs`), so audit storage is bounded and, without adjustment, the oldest records are eventually overwritten. Size these settings and the disk-full and low-space actions to your site's audit-retention policy so that audit data is not lost.

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.

Impact:

Log rotation at the configured `max_log_file` size keeps individual audit log files bounded. Sizing it too small rotates logs frequently (and, combined with a rotate/overwrite action, risks losing older events); sizing it too large can consume disk space. Size per site policy alongside the disk-full and low-space handling.

Audit:

Run the following command to verify the `max_log_file` parameter is configured in accordance 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 :: AU-4
3. AUDITD.CONF(5)

Additional Information:

The `max_log_file` parameter is measured in megabytes.

Other methods of log rotation may be appropriate based on site policy. One example is time-based rotation strategies which don't have native support in `auditd` configurations. Manual audit of custom configurations should be evaluated for effectiveness and completeness.

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	8.3 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.		●	●

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.

Impact:

With `max_log_file_action = keep_logs`, audit logs are rotated but never automatically deleted, so audit log disk usage grows until logs are manually managed. Ensure adequate audit log storage is configured, that the disk-full and low-space handling (see "Ensure system is disabled when audit logs are full" and "Ensure system warns when audit logs are low on space") are in place, and monitor consumption.

Audit:

Run the following command to verify the `max_log_file_action` parameter is set to `keep_logs`:

```
# grep -Pi -- '^\h*max_log_file_action\h*=\h*keep_logs\b'
/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
```

Default Value:

`max_log_file_action = ROTATE`

References:

1. GRID: NIX-00000295
2. NIST SP 800-53 :: AU-11
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.		●	●

6.2.2.3 Ensure system is disabled when audit logs are full (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

The `disk_full_action` parameter sets the action `auditd` takes when the partition holding the audit log files has no free space. The `disk_error_action` parameter sets the action `auditd` takes when an error is detected on that partition.

Rationale:

In high security contexts, the risk of detecting unauthorized access or nonrepudiation exceeds the benefit of the system's availability.

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:

Both of the following must be true.

1. Run the following command to verify `disk_full_action` is set to `halt` or `single`:

```
# grep -Pi -- '^\h*disk_full_action\h*=\h*(halt|single)\b'
/etc/audit/auditd.conf

disk_full_action = <halt|single>
```

2. Run the following command to verify `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:

```
disk_full_action = halt
disk_error_action = halt
```

Default Value:

`disk_full_action = SUSPEND`

`disk_error_action = SYSLOG`

References:

1. GRID: NIX-00000338
2. AUDITD.CONF(5)
3. https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html/security_hardening/auditing-the-system_security-hardening#configuring-auditd-for-a-secure-environment_auditing-the-system
4. NIST SP 800-53 :: AU-5 b
5. NIST SP 800-53A :: AU-5.1 (iv)
6. RHEL 8 STIG Vul ID: V-230390
7. RHEL 8 STIG Rule ID: SV-230390r627750_rule
8. RHEL 8 STIG Vul ID: V-230392
9. RHEL 8 STIG Rule ID: SV-230392r627750

Additional Information:

`disk_full_action` valid values are `ignore`, `syslog`, `rotate`, `exec`, `suspend`, `single`, and `halt`:

- `ignore` - the audit daemon does nothing
- `syslog` - the audit daemon issues a warning to syslog
- `rotate` - the audit daemon rotates the logs, discarding the oldest to free space
- `exec /path-to-script` - runs the script (no parameters); the script must resume auditd logging when done
- `suspend` - the audit daemon stops writing records to disk
- `single` - the audit daemon puts the system in single user mode
- `halt` - the audit daemon shuts the system down

`disk_error_action` valid values are `ignore`, `syslog`, `exec`, `suspend`, `single`, and `halt` (same meanings; `syslog` issues no more than 5 consecutive warnings).

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.		●	●

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 `space_left_action` parameter sets the action `auditd` takes when free space on the partition holding the audit log files drops below the `space_left` threshold. The `admin_space_left_action` parameter sets the action `auditd` takes when free space drops below the lower `admin_space_left` threshold.

Rationale:

In high security contexts, the risk of detecting unauthorized access or nonrepudiation exceeds the benefit of the system's availability.

Impact:

If `space_left_action` is set to `email`, a Mail Transfer Agent must be available to deliver the low-space warning. If `space_left_action` or `admin_space_left_action` is set to `single`, `auditd` puts the system in single user mode when the corresponding threshold is reached; if set to `halt`, `auditd` shuts the system down. Choose actions consistent with your availability requirements.

Audit:

Both of the following must be true.

1. Run the following command to verify `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

space_left_action = email
```

2. Run the following command to verify `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

admin_space_left_action = single
```

Remediation:

1. Set the `space_left_action` parameter in `/etc/audit/auditd.conf` to `email`, `exec`, `single`, or `halt`:

```
space_left_action = email
```

2. Set the `admin_space_left_action` parameter in `/etc/audit/auditd.conf` to `single` or `halt`:

```
admin_space_left_action = single
```

Note: A Mail Transfer Agent (MTA) must be installed and configured to use `space_left_action = email`.

Default Value:

`space_left_action = SYSLOG`

`admin_space_left_action = SUSPEND`

References:

1. GRID: NIX-00000339
2. NIST SP 800-53 :: AU-5
3. AUDITD.CONF(5)
4. https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html/security_hardening/auditing-the-system_security-hardening#configuring-auditd-for-a-secure-environment_auditing-the-system

Additional Information:

`space_left_action` valid values are `ignore`, `syslog`, `rotate`, `email`, `exec`, `suspend`, `single`, and `halt`. `admin_space_left_action` accepts the same values. `email` sends a warning to the account named by `action_mail_acct` and requires a configured Mail Transfer Agent (MTA); the other values behave as for `disk_full_action`.

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.		●	●

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).

  These watch rules generally follow the 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: `-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.

General guidance for auditd rules (applies to all recommendations in this section)

Rule format. Each rule begins with `-a always,exit`. A watch is anchored by a `-F path=<file>` or `-F dir=<directory>` target together with a `-F perm=<perms>` permission set (e.g., `wa`; broader sets such as `rwxa` are also acceptable). A `-k <name>` / `-F key=<name>` label is recommended (any name may be used). `-S all` places no restriction on the system call - for a path/dir-anchored watch the rule triggers on any system call that accesses the target with the listed permissions; `auditctl` adds `-S all` automatically if it is omitted from a watch rule.

Loading and immutability. Rules placed in `/etc/audit/rules.d/*.rules` are merged and loaded with `augenrules --load`. If the configuration is immutable (`-e 2`), new rules are not loaded into the running configuration and a reboot is required to apply them (and `augenrules --load` will not warn that they were not loaded):

```
# augenrules --load
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

Architecture (`arch=`). Supplying `-F arch=` is recommended for performance (omitting it evaluates every system call). On EL9-family x86_64 systems both the 64-bit and 32-bit Application Binary Interfaces (ABIs) are present by default, so the rules in this benchmark target **both** (`-F arch=b64` and `-F arch=b32`); a separate rule is included for each ABI so that 32-bit (compat) system calls are audited as well as 64-bit ones. On non-x86 architectures (e.g., aarch64, ppc64le, s390x) the system-call set and the correct `-F arch=` token differ; verify the appropriate `-F arch=` filter for the platform and confirm the rules load with `augenrules --load`.

Deprecated `-w` (watch) form. The legacy `-w <path> -p <perms>` form is deprecated; it still functions but is expected to be removed in a future release. Either the legacy `-w` form or the modern `-a always,exit ... -F path=<path>`/`-F dir=<path>` ... `-F perm=<perms>` form is acceptable, and **both are a passing state** for this section's automated checks - a watch already written in the legacy `-w` form is reported as **passing and does not need to be changed**, though updating it to the modern form is recommended. The Remediation Procedures create rules in the modern form only.

`auid>=` / UID_MIN. Where a rule includes `-F auid>=1000`, use the system's `UID_MIN` value from `/etc/login.defs` (1000 on a default install).

Watch-target existence. An auditd rule whose `-F path=`/`-F dir=` (or legacy `-w`) target does not exist will fail to load. A watch rule is therefore added only when its target is present; if the target does not exist on a given host, the recommendation does not apply to that host until the target is created (re-run the remediation afterward).

Verifying audit rules. Rules are written to `/etc/audit/rules.d/*.rules` in the form shown in each recommendation's Remediation. Multiple system calls may be placed on a single rule line (preferred for performance - see `man 7 audit.rules`) or split across several; the groupings are equivalent. When the rules are loaded, the kernel normalizes some fields, so `auditctl -l` output differs from the on-disk text - for example `-F auid!=unset` is shown as `-F auid!=-1`, `-k <key>` as `-F key=<key>`, and `-C euid!=uid` as `-C uid!=euid`. The order of rules, of system calls within a `-S` clause, and of fields may also differ. A rule is satisfied when an **equivalent** rule is present, regardless of these differences.

Self-determining remediation. Each remediation derives its rules from the running kernel rather than hard-coding a fixed set, so it does not add rules the kernel cannot use. For each ABI, a rule is included for a given system call only when that call exists for that ABI (resolved with `ausyscall <b64|b32> <name>`), so version-specific calls (for example, `fchmodat2` on newer kernels, or `create_module`/`query_module` on older ones) are included or omitted automatically; watch rules are added only when their path or directory exists. This keeps each recommendation portable across kernel versions and across both ABIs without manual edits.

Audit log volume. Enabling the audit rules in this section increases the number of records `auditd` writes and therefore audit log disk usage. Rules covering high-frequency system calls (for example `chmod`, `chown`, `open`/`openat`, `execve`, `unlink`) generate the most records and can have a minor performance impact under heavy load; rules that watch configuration files or specific commands generate records only when those targets are modified or used. Size audit log storage and configure rotation/retention accordingly (see the "Configure Data Retention" recommendations).

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:

```
#!/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
```

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.

Audit:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- 'scope' /etc/audit/rules.d/*.rules
```

Verify a watch rule covering each target below is present — in either the modern `-F path=`/`-F dir=` form shown (preferred) or the legacy `-w <path> -p <perms>` form; both are a passing state:

```
-a always,exit -S all -F path=/etc/sudoers -F perm=wa -F key=scope
-a always,exit -S all -F dir=/etc/sudoers.d -F perm=wa -F key=scope
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- 'scope'
```

Verify the same watch rules are loaded (either form is accepted):

```
-a always,exit -S all -F path=/etc/sudoers -F perm=wa -F key=scope
-a always,exit -S all -F dir=/etc/sudoers.d -F perm=wa -F key=scope
```

Remediation:

1. Add the audit rule(s) that monitor changes to `/etc/sudoers` and `/etc/sudoers.d`. The procedure adds each rule only for targets that exist, so `augenrules --load` does not error. Run:

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

{
   [ -e /etc/sudoers ] && printf '%s\n' "-a always,exit -S all -F
path=/etc/sudoers -F perm=wa -F key=scope"
   [ -e /etc/sudoers.d ] && printf '%s\n' "-a always,exit -S all -F
dir=/etc/sudoers.d -F perm=wa -F key=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
```

Default Value:

By default, no audit rule is present to monitor changes to `/etc/sudoers` and `/etc/sudoers.d`.

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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	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.		●	●

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.

Audit:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- 'execve' /etc/audit/rules.d/*.rules
```

Verify a rule equivalent to each line below is present:

```
-a always,exit -F arch=b64 -C euid!=uid -F auid!=unset -S execve,execveat -k
user_emulation
-a always,exit -F arch=b32 -C euid!=uid -F auid!=unset -S execve,execveat -k
user_emulation
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- 'execve'
```

Verify the same rules are loaded; `auditctl -l` displays them in the normalized form below:

```
-a always,exit -F arch=b64 -S execve,execveat -C uid!=euid -F auid!=-1 -F
key=user_emulation
-a always,exit -F arch=b32 -S execve,execveat -C uid!=euid -F auid!=-1 -F
key=user_emulation
```

Remediation:

1. Add the audit rule(s) for the `execve` and `execveat` system calls. The procedure adds a rule for each audit ABI (b64/b32) the running kernel supports and skips any unsupported ABI or system call, so `augenrules --load` does not error. Run:

```
for ARCH in b64 b32; do
   CALLS=""; for C in execve execveat; do ausyscall "$ARCH" "$C" >/dev/null
2>&1 && CALLS="${CALLS:+$CALLS,}$C"; done
   [ -n "$CALLS" ] && printf '%s\n' "-a always,exit -F arch=$ARCH -C
euid!=uid -F auid!=unset -S $CALLS -k user_emulation"
done >> /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
```

Default Value:

By default, no audit rule is present to record the `execve` system call.

References:

1. GRID: NIX-00000286
2. NIST SP 800-53 Revision 5 :: AC-6 (8)
3. RHEL 8 STIG Vul ID: V-230386
4. RHEL 8 STIG Rule ID: SV-230386r854037

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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	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.		●	●

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:

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 the rule is present in a rules file:

```
# grep -Ps -- 'sudo.log' /etc/audit/rules.d/*.rules
```

Verify a watch rule covering the target below is present — in either the modern `-F path=`/`-F dir=` form shown (preferred) or the legacy `-w <path> -p <perms>` form; both are a passing state:

```
-a always,exit -S all -F path=/var/log/sudo.log -F perm=wa -F
key=sudo_log_file
```

3. Run the following command to verify the rule is loaded in the running configuration:

```
# auditctl -l | grep -Ps -- 'sudo.log'
```

Verify the same watch rule is loaded (either form is accepted):

```
-a always,exit -S all -F path=/var/log/sudo.log -F perm=wa -F
key=sudo_log_file
```

Remediation:

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. The rule is added only if the configured log file exists.

Example:

```
[ -e /var/log/sudo.log ] && printf '%s\n' "-a always,exit -S all -F
path=/var/log/sudo.log -F perm=wa -F key=sudo_log_file" >>
/etc/audit/rules.d/50-sudo.rules
```

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
```

Default Value:

By default, no audit rule is present to monitor the configured sudo log file.

References:

1. GRID: NIX-00000320

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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	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.		●	●

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
- `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:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- '(adjtimex|settimeofday|clock_settime|/etc/localtime)'
/etc/audit/rules.d/*.rules
```

Verify a rule equivalent to each line below is present:

```
-a always,exit -F arch=b64 -S adjtimex,settimeofday -k time-change
-a always,exit -F arch=b32 -S adjtimex,settimeofday -k time-change
-a always,exit -F arch=b64 -S clock_settime -F a0=0x0 -k time-change
-a always,exit -F arch=b32 -S clock_settime -F a0=0x0 -k time-change
-a always,exit -S all -F path=/etc/localtime -F perm=wa -k time-change
```

The `/etc/localtime` watch (last line) may be written in either the modern `-F path=` form shown or the legacy `-w /etc/localtime -p wa` form; both are a passing state (the system-call rules above have no `-w` equivalent).

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps --
'(adjtimex|settimeofday|clock_settime|/etc/localtime)'
```

Verify the same rules are loaded; `auditctl -l` displays them in the normalized form below:

```
-a always,exit -F arch=b64 -S adjtimex,settimeofday -F key=time-change
-a always,exit -F arch=b32 -S adjtimex,settimeofday -F key=time-change
-a always,exit -F arch=b64 -S clock_settime -F a0=0x0 -F key=time-change
-a always,exit -F arch=b32 -S clock_settime -F a0=0x0 -F key=time-change
-a always,exit -S all -F path=/etc/localtime -F perm=wa -F key=time-change
```

Remediation:

1. Add the audit rule(s) for the `adjtimex`, `settimeofday`, and `clock_settime` system calls and changes to `/etc/localtime`. The procedure adds a rule for each audit ABI (b64/b32) the running kernel supports, skips any unsupported ABI or system call, and adds the watch only when its target exists, so `augenrules --load` does not error. Run:

```
for ARCH in b64 b32; do
   CALLS=""; for C in adjtimex settimeofday; do ausyscall "$ARCH" "$C"
>/dev/null 2>&1 && CALLS="${CALLS:+$CALLS,}$C"; done
   [ -n "$CALLS" ] && printf '%s\n' "-a always,exit -F arch=$ARCH -S $CALLS -
k time-change"
   CALLS=""; for C in clock_settime; do ausyscall "$ARCH" "$C" >/dev/null
2>&1 && CALLS="${CALLS:+$CALLS,}$C"; done
   [ -n "$CALLS" ] && printf '%s\n' "-a always,exit -F arch=$ARCH -S $CALLS -
F a0=0x0 -k time-change"
done >> /etc/audit/rules.d/50-time-change.rules
[ -e /etc/localtime ] && printf '%s\n' "-a always,exit -S all -F
path=/etc/localtime -F perm=wa -F key=time-change" >> /etc/audit/rules.d/50-
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
```

Default Value:

By default, no audit rule is present to record the `adjtimex`, `settimeofday`, and `clock_settime` system calls and changes to `/etc/localtime`.

References:

1. GRID: NIX-00000318
2. NIST SP 800-53 Rev. 5: AU-3, CM-6

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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	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.		●	●

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:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- 'sethostname' /etc/audit/rules.d/*.rules
```

Verify a rule equivalent to each line below is present:

```
-a always,exit -F arch=b64 -S sethostname,setdomainname -k system-locale
-a always,exit -F arch=b32 -S sethostname,setdomainname -k system-locale
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- 'sethostname'
```

Verify the same rules are loaded; `auditctl -l` displays them in the normalized form below:

```
-a always,exit -F arch=b64 -S sethostname,setdomainname -F key=system-locale
-a always,exit -F arch=b32 -S sethostname,setdomainname -F key=system-locale
```

Remediation:

1. Add the audit rule(s) for the `sethostname` and `setdomainname` system calls. The procedure adds a rule for each audit ABI (b64/b32) the running kernel supports and skips any unsupported ABI or system call, so `augenrules --load` does not error. Run:

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

{
   for ARCH in b64 b32; do
      CALLS=""; for C in sethostname setdomainname; do ausyscall "$ARCH" "$C"
>/dev/null 2>&1 && CALLS="${CALLS:+$CALLS,}$C"; done
      [ -n "$CALLS" ] && printf '%s\n' "-a always,exit -F arch=$ARCH -S
$CALLS -k system-locale"
   done >> /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
```

Default Value:

By default, no audit rule is present to record the `sethostname` and `setdomainname` system calls.

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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	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.		●	●

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:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- '/etc/issue' /etc/audit/rules.d/*.rules
```

Verify a watch rule covering each target below is present — in either the modern `-F path=`/`-F dir=` form shown (preferred) or the legacy `-w <path> -p <perms>` form; both are a passing state:

```
-a always,exit -S all -F path=/etc/issue -F perm=wa -F key=system-locale
-a always,exit -S all -F path=/etc/issue.net -F perm=wa -F key=system-locale
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- '/etc/issue'
```

Verify the same watch rules are loaded (either form is accepted):

```
-a always,exit -S all -F path=/etc/issue -F perm=wa -F key=system-locale
-a always,exit -S all -F path=/etc/issue.net -F perm=wa -F key=system-locale
```

Remediation:

1. Add the audit rule(s) that monitor changes to `/etc/issue` and `/etc/issue.net`. The procedure adds each rule only for targets that exist, so `augenrules --load` does not error. Run:

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

{
   [ -e /etc/issue ] && printf '%s\n' "-a always,exit -S all -F
path=/etc/issue -F perm=wa -F key=system-locale"
   [ -e /etc/issue.net ] && printf '%s\n' "-a always,exit -S all -F
path=/etc/issue.net -F perm=wa -F key=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
```

Default Value:

By default, no audit rule is present to monitor changes to `/etc/issue` and `/etc/issue.net`.

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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	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.		●	●

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:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- '/etc/host' /etc/audit/rules.d/*.rules
```

Verify a watch rule covering each target below is present — in either the modern `-F path=`/`-F dir=` form shown (preferred) or the legacy `-w <path> -p <perms>` form; both are a passing state:

```
-a always,exit -S all -F path=/etc/hosts -F perm=wa -F key=system-locale
-a always,exit -S all -F path=/etc/hostname -F perm=wa -F key=system-locale
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- '/etc/host'
```

Verify the same watch rules are loaded (either form is accepted):

```
-a always,exit -S all -F path=/etc/hosts -F perm=wa -F key=system-locale
-a always,exit -S all -F path=/etc/hostname -F perm=wa -F key=system-locale
```

Remediation:

1. Add the audit rule(s) that monitor changes to `/etc/hosts` and `/etc/hostname`. The procedure adds each rule only for targets that exist, so `augenrules --load` does not error. Run:

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

{
   [ -e /etc/hosts ] && printf '%s\n' "-a always,exit -S all -F
path=/etc/hosts -F perm=wa -F key=system-locale"
   [ -e /etc/hostname ] && printf '%s\n' "-a always,exit -S all -F
path=/etc/hostname -F perm=wa -F key=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
```

Default Value:

By default, no audit rule is present to monitor changes to `/etc/hosts` and `/etc/hostname`.

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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	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.		●	●

6.2.3.8 Ensure events that modify /etc/sysconfig/network and /etc/NetworkManager/system-connections/ 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/sysconfig/network` - additional information that is valid to all network interfaces
- `/etc/NetworkManager/system-connections/` - directory containing user-created, persistent network configuration files in a keyfile format.

Rationale:

Monitoring /etc/sysconfig/network is important as it can show if network interfaces or scripts are being modified in a way that can lead to the machine becoming unavailable or compromised. All audit records should have a relevant tag associated with them.

Audit:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- '(sysconfig/network|system-connections)'
/etc/audit/rules.d/*.rules
```

Verify a watch rule covering each target below is present — in either the modern `-F path=`/`-F dir=` form shown (preferred) or the legacy `-w <path> -p <perms>` form; both are a passing state:

```
-a always,exit -S all -F path=/etc/sysconfig/network -F perm=wa -F
key=system-locale
-a always,exit -S all -F dir=/etc/NetworkManager/system-connections -F
perm=wa -F key=system-locale
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- '(sysconfig/network|system-connections)'
```

Verify the same watch rules are loaded (either form is accepted):

```
-a always,exit -S all -F path=/etc/sysconfig/network -F perm=wa -F
key=system-locale
-a always,exit -S all -F dir=/etc/NetworkManager/system-connections -F
perm=wa -F key=system-locale
```

Remediation:

1. Add the audit rule(s) that monitor changes to `/etc/sysconfig/network` and `/etc/NetworkManager/system-connections`. The procedure adds each rule only for targets that exist, so `augenrules --load` does not error. Run:

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

{
   [ -e /etc/sysconfig/network ] && printf '%s\n' "-a always,exit -S all -F
path=/etc/sysconfig/network -F perm=wa -F key=system-locale"
   [ -d /etc/NetworkManager/system-connections ] && printf '%s\n' "-a
always,exit -S all -F dir=/etc/NetworkManager/system-connections -F perm=wa -
F key=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
```

Default Value:

By default, no audit rule is present to monitor changes to `/etc/sysconfig/network` and `/etc/NetworkManager/system-connections`.

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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	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.		●	●

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:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- '/etc/NetworkManager' /etc/audit/rules.d/*.rules
```

Verify a watch rule covering the target below is present — in either the modern `-F path=`/`-F dir=` form shown (preferred) or the legacy `-w <path> -p <perms>` form; both are a passing state:

```
-a always,exit -S all -F dir=/etc/NetworkManager -F perm=wa -F key=system-
locale
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- '/etc/NetworkManager'
```

Verify the same watch rule is loaded (either form is accepted):

```
-a always,exit -S all -F dir=/etc/NetworkManager -F perm=wa -F key=system-
locale
```

Remediation:

1. Add the audit rule(s) that monitor changes to `/etc/NetworkManager`. The procedure adds each rule only for targets that exist, so `augenrules --load` does not error. Run:

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

{
   [ -d /etc/NetworkManager ] && printf '%s\n' "-a always,exit -S all -F
dir=/etc/NetworkManager -F perm=wa -F key=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
```

Default Value:

By default, no audit rule is present to monitor changes to `/etc/NetworkManager`.

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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	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.		●	●

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:

Privileged programs are those with the setuid and/or setgid bit set on execution. For each such program on a device-backed filesystem (excluding `noexec`/`nosuid` mounts and the ephemeral container/snap paths noted below), verify an auditd rule records its execution by unprivileged users, both on disk and in the running configuration.

On disk configuration

Run the following script. Nothing should be returned (any output names a privileged program lacking a compliant on-disk rule):

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

{
   UID_MIN="$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)"
   for PARTITION in $(findmnt -n -l -k -it $(awk '/nodev/ { print $2 }'
/proc/filesystems | paste -sd,) | grep -Pv "noexec|nosuid" | awk '{print
$1}'); do
      while IFS= read -r FILE; do
         ok=""
         while IFS= read -r L; do
            grep -Pq -- '-a\s+(always,exit|exit,always)' <<< "$L" \
            && grep -Pq -- '-F\s+perm=x' <<< "$L" \
            && grep -Pq -- "-F\s+auid>=${UID_MIN}" <<< "$L" \
            && grep -Pq -- '-F\s+auid!=(unset|-1|4294967295)' <<< "$L" &&
ok="Y"
         done < <(grep -h -- "path=${FILE}" /etc/audit/rules.d/*.rules
2>/dev/null)
         [ -z "$ok" ] && printf 'MISSING or incomplete on-disk rule: %s\n'
"$FILE"
      done < <(find "${PARTITION}" -xdev \( -path "*/containers/storage/*" -o
-path "*/containerd/*" -o -path "*/kubelet/*" -o -path "/snap/*" \) -prune -o
-perm /6000 -type f -print)
   done
}
```

Running configuration

Run the following script. Nothing should be returned (any output names a privileged program not covered by a loaded rule):

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

{
   UID_MIN="$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)"
   RUNNING="$(auditctl -l)"
   for PARTITION in $(findmnt -n -l -k -it $(awk '/nodev/ { print $2 }'
/proc/filesystems | paste -sd,) | grep -Pv "noexec|nosuid" | awk '{print
$1}'); do
      while IFS= read -r FILE; do
         ok=""
         while IFS= read -r L; do
            grep -Pq -- '-a\s+(always,exit|exit,always)' <<< "$L" \
            && grep -Pq -- '-F\s+perm=x' <<< "$L" \
            && grep -Pq -- "-F\s+auid>=${UID_MIN}" <<< "$L" \
            && grep -Pq -- '-F\s+auid!=(unset|-1|4294967295)' <<< "$L" &&
ok="Y"
         done < <(printf '%s\n' "$RUNNING" | grep -- "path=${FILE}")
         [ -z "$ok" ] && printf 'MISSING or incomplete loaded rule: %s\n'
"$FILE"
      done < <(find "${PARTITION}" -xdev \( -path "*/containers/storage/*" -o
-path "*/containerd/*" -o -path "*/kubelet/*" -o -path "/snap/*" \) -prune -o
-perm /6000 -type f -print)
   done
}
```

The running configuration normalizes some fields (for example `-F auid!=unset` is shown as `-F auid!=-1`, and `-k key` as `-F key=key`); the checks above accept these equivalent forms.

Excluded paths. `noexec`/`nosuid` mounts confer no execution or privilege and are out of scope. Ephemeral, runtime-managed container and snap layer paths (`*/containers/storage/*`, `*/containerd/*`, `*/kubelet/*`, `/snap/*`) are also excluded: their setuid files live in versioned/overlay paths that churn and execute inside container namespaces, so a host `-F path=` rule cannot usefully audit them. The control targets unprivileged use of host privileged commands.

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:

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

{
  UID_MIN=$(awk '/^\s*UID_MIN/{print $2}' /etc/login.defs)
  AUDIT_RULE_FILE="/etc/audit/rules.d/50-privileged.rules"
  # Emit a rule only for the audit ABIs the running kernel supports
  ARCHES=""; for ARCH in b64 b32; do ausyscall "$ARCH" execve >/dev/null 2>&1
&& ARCHES="${ARCHES:+$ARCHES }$ARCH"; done
  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
    while IFS= read -r FILE; do
      for ARCH in $ARCHES; do
        NEW_DATA+=("-a always,exit -F arch=$ARCH -S all -F path=${FILE} -F
perm=x -F auid>=${UID_MIN} -F auid!=unset -k privileged")
      done
    done < <(find "${PARTITION}" -xdev \( -path "*/containers/storage/*" -o -
path "*/containerd/*" -o -path "*/kubelet/*" -o -path "/snap/*" \) -prune -o
-perm /6000 -type f -print)
  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.

Excluded paths. `noexec`/`nosuid` mounts, and the ephemeral container/snap layer paths pruned above (`*/containers/storage/*`, `*/containerd/*`, `*/kubelet/*`, `/snap/*`), are intentionally out of scope: those setuid files live in versioned/overlay paths that churn and execute inside container namespaces, so a host `-F path=` rule cannot usefully audit them. The recommendation targets unprivileged use of host privileged commands.

Default Value:

By default, no audit rules are present to monitor the use of privileged (setuid/setgid) commands.

References:

1. GRID: NIX-00000346
2. NIST SP 800-53 Rev. 5: AU-3, AU-3(1)

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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.	●	●	●

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.

Audit:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- '(EACCES|EPERM)' /etc/audit/rules.d/*.rules
```

Verify a rule equivalent to each line below is present, one per audit ABI the kernel supports:

```
-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=b32 -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
-a always,exit -F arch=b32 -S creat,open,openat,truncate,ftruncate -F exit=-
EPERM -F auid>=1000 -F auid!=unset -k access
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- '(EACCES|EPERM)'
```

Verify the same rules are loaded; `auditctl -l` displays them in the normalized form below:

```
-a always,exit -F arch=b64 -S creat,open,openat,truncate,ftruncate -F exit=-
EACCES -F auid>=1000 -F auid!=-1 -F key=access
-a always,exit -F arch=b32 -S creat,open,openat,truncate,ftruncate -F exit=-
EACCES -F auid>=1000 -F auid!=-1 -F key=access
-a always,exit -F arch=b64 -S creat,open,openat,truncate,ftruncate -F exit=-
EPERM -F auid>=1000 -F auid!=-1 -F key=access
-a always,exit -F arch=b32 -S creat,open,openat,truncate,ftruncate -F exit=-
EPERM -F auid>=1000 -F auid!=-1 -F key=access
```

Remediation:

1. Add the audit rule(s) for the `creat`, `open`, `openat`, `truncate`, and `ftruncate` system calls. The procedure adds a rule for each audit ABI (b64/b32) the running kernel supports and skips any unsupported ABI or system call, so `augenrules --load` does not error. Run:

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

{
   UID_MIN="$(awk '/^[[:space:]]*UID_MIN/{print $2}' /etc/login.defs)"
   for ARCH in b64 b32; do
      CALLS=""; for C in creat open openat truncate ftruncate; do ausyscall
"$ARCH" "$C" >/dev/null 2>&1 && CALLS="${CALLS:+$CALLS,}$C"; done
      [ -n "$CALLS" ] && printf '%s\n' "-a always,exit -F arch=$ARCH -S
$CALLS -F exit=-EACCES -F auid>=$UID_MIN -F auid!=unset -k access"
      [ -n "$CALLS" ] && printf '%s\n' "-a always,exit -F arch=$ARCH -S
$CALLS -F exit=-EPERM -F auid>=$UID_MIN -F auid!=unset -k access"
   done >> /etc/audit/rules.d/50-access.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
```

Default Value:

By default, no audit rule is present to record the `creat`, `open`, `openat`, `truncate`, and `ftruncate` system calls.

References:

1. GRID: NIX-00000345
2. NIST SP 800-53 :: AU-12 a
3. NIST SP 800-53A :: AU-12.1 (ii)
4. RHEL 8 STIG Vul ID: V-230449
5. RHEL 8 STIG Rule ID: SV-230449r810455

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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	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).			●

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.

Audit:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- '/etc/group' /etc/audit/rules.d/*.rules
```

Verify a watch rule covering the target below is present — in either the modern `-F path=`/`-F dir=` form shown (preferred) or the legacy `-w <path> -p <perms>` form; both are a passing state:

```
-a always,exit -S all -F path=/etc/group -F perm=wa -F key=identity
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- '/etc/group'
```

Verify the same watch rule is loaded (either form is accepted):

```
-a always,exit -S all -F path=/etc/group -F perm=wa -F key=identity
```

Remediation:

1. Add the audit rule(s) that monitor changes to `/etc/group`. The procedure adds each rule only for targets that exist, so `augenrules --load` does not error. Run:

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

{
   [ -e /etc/group ] && printf '%s\n' "-a always,exit -S all -F
path=/etc/group -F perm=wa -F key=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
```

Default Value:

By default, no audit rule is present to monitor changes to `/etc/group`.

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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	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.		●	●

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.

Audit:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- '/etc/passwd' /etc/audit/rules.d/*.rules
```

Verify a watch rule covering the target below is present — in either the modern `-F path=`/`-F dir=` form shown (preferred) or the legacy `-w <path> -p <perms>` form; both are a passing state:

```
-a always,exit -S all -F path=/etc/passwd -F perm=wa -F key=identity
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- '/etc/passwd'
```

Verify the same watch rule is loaded (either form is accepted):

```
-a always,exit -S all -F path=/etc/passwd -F perm=wa -F key=identity
```

Remediation:

1. Add the audit rule(s) that monitor changes to `/etc/passwd`. The procedure adds each rule only for targets that exist, so `augenrules --load` does not error. Run:

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

{
   [ -e /etc/passwd ] && printf '%s\n' "-a always,exit -S all -F
path=/etc/passwd -F perm=wa -F key=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
```

Default Value:

By default, no audit rule is present to monitor changes to `/etc/passwd`.

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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	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.		●	●

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.

Audit:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- 'shadow' /etc/audit/rules.d/*.rules
```

Verify a watch rule covering each target below is present — in either the modern `-F path=`/`-F dir=` form shown (preferred) or the legacy `-w <path> -p <perms>` form; both are a passing state:

```
-a always,exit -S all -F path=/etc/shadow -F perm=wa -F key=identity
-a always,exit -S all -F path=/etc/gshadow -F perm=wa -F key=identity
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- 'shadow'
```

Verify the same watch rules are loaded (either form is accepted):

```
-a always,exit -S all -F path=/etc/shadow -F perm=wa -F key=identity
-a always,exit -S all -F path=/etc/gshadow -F perm=wa -F key=identity
```

Remediation:

1. Add the audit rule(s) that monitor changes to `/etc/gshadow` and `/etc/shadow`. The procedure adds each rule only for targets that exist, so `augenrules --load` does not error. Run:

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

{
   [ -e /etc/gshadow ] && printf '%s\n' "-a always,exit -S all -F
path=/etc/gshadow -F perm=wa -F key=identity"
   [ -e /etc/shadow ] && printf '%s\n' "-a always,exit -S all -F
path=/etc/shadow -F perm=wa -F key=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
```

Default Value:

By default, no audit rule is present to monitor changes to `/etc/gshadow` and `/etc/shadow`.

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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	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.		●	●

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.

Audit:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- 'opasswd' /etc/audit/rules.d/*.rules
```

Verify a watch rule covering the target below is present — in either the modern `-F path=`/`-F dir=` form shown (preferred) or the legacy `-w <path> -p <perms>` form; both are a passing state:

```
-a always,exit -S all -F path=/etc/security/opasswd -F perm=wa -F
key=identity
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- 'opasswd'
```

Verify the same watch rule is loaded (either form is accepted):

```
-a always,exit -S all -F path=/etc/security/opasswd -F perm=wa -F
key=identity
```

Remediation:

1. Add the audit rule(s) that monitor changes to `/etc/security/opasswd`. The procedure adds each rule only for targets that exist, so `augenrules --load` does not error. Run:

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

{
   [ -e /etc/security/opasswd ] && printf '%s\n' "-a always,exit -S all -F
path=/etc/security/opasswd -F perm=wa -F key=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
```

Default Value:

By default, no audit rule is present to monitor changes to `/etc/security/opasswd`.

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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	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.		●	●

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.

Audit:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- 'nsswitch' /etc/audit/rules.d/*.rules
```

Verify a watch rule covering the target below is present — in either the modern `-F path=`/`-F dir=` form shown (preferred) or the legacy `-w <path> -p <perms>` form; both are a passing state:

```
-a always,exit -S all -F path=/etc/nsswitch.conf -F perm=wa -F key=identity
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- 'nsswitch'
```

Verify the same watch rule is loaded (either form is accepted):

```
-a always,exit -S all -F path=/etc/nsswitch.conf -F perm=wa -F key=identity
```

Remediation:

1. Add the audit rule(s) that monitor changes to `/etc/nsswitch.conf`. The procedure adds each rule only for targets that exist, so `augenrules --load` does not error. Run:

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

{
   [ -e /etc/nsswitch.conf ] && printf '%s\n' "-a always,exit -S all -F
path=/etc/nsswitch.conf -F perm=wa -F key=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
```

Default Value:

By default, no audit rule is present to monitor changes to `/etc/nsswitch.conf`.

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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	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.		●	●

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.

Audit:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- 'pam' /etc/audit/rules.d/*.rules
```

Verify a watch rule covering each target below is present — in either the modern `-F path=`/`-F dir=` form shown (preferred) or the legacy `-w <path> -p <perms>` form; both are a passing state:

```
-a always,exit -S all -F path=/etc/pam.conf -F perm=wa -F key=identity
-a always,exit -S all -F dir=/etc/pam.d -F perm=wa -F key=identity
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- 'pam'
```

Verify the same watch rules are loaded (either form is accepted):

```
-a always,exit -S all -F path=/etc/pam.conf -F perm=wa -F key=identity
-a always,exit -S all -F dir=/etc/pam.d -F perm=wa -F key=identity
```

Remediation:

1. Add the audit rule(s) that monitor changes to `/etc/pam.conf` and `/etc/pam.d`. The procedure adds each rule only for targets that exist, so `augenrules --load` does not error. Run:

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

{
   [ -e /etc/pam.conf ] && printf '%s\n' "-a always,exit -S all -F
path=/etc/pam.conf -F perm=wa -F key=identity"
   [ -d /etc/pam.d ] && printf '%s\n' "-a always,exit -S all -F
dir=/etc/pam.d -F perm=wa -F key=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
```

Default Value:

By default, no audit rule is present to monitor changes to `/etc/pam.conf` and `/etc/pam.d`.

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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	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.		●	●

6.2.3.18 Ensure discretionary access control permission modification events chmod,fchmod,fchmodat 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

In all cases, an audit record will only be written for non-system user ids and will ignore Daemon events. All audit records will be tagged with the identifier "perm_mod."

Rationale:

Monitoring for changes in file attributes could alert a system administrator to activity that could indicate intruder activity or policy violation.

Audit:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- 'chmod' /etc/audit/rules.d/*.rules
```

Verify a rule equivalent to each line below is present:

```
-a always,exit -F arch=b64 -S chmod,fchmod,fchmodat -F auid>=1000 -F
auid!=unset -F key=perm_mod
-a always,exit -F arch=b32 -S chmod,fchmod,fchmodat -F auid>=1000 -F
auid!=unset -F key=perm_mod
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- 'chmod'
```

Verify the same rules are loaded; `auditctl -l` displays them in the normalized form below:

```
-a always,exit -F arch=b64 -S chmod,fchmod,fchmodat -F auid>=1000 -F auid!=-1
-F key=perm_mod
-a always,exit -F arch=b32 -S chmod,fchmod,fchmodat -F auid>=1000 -F auid!=-1
-F key=perm_mod
```

Remediation:

1. Add the audit rule(s) for the `chmod`, `fchmod`, and `fchmodat` system calls. The procedure adds a rule for each audit ABI (b64/b32) the running kernel supports and skips any unsupported ABI or system call, so `augenrules --load` does not error. Run:

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

{
   UID_MIN="$(awk '/^[[:space:]]*UID_MIN/{print $2}' /etc/login.defs)"
   for ARCH in b64 b32; do
      CALLS=""; for C in chmod fchmod fchmodat; do ausyscall "$ARCH" "$C"
>/dev/null 2>&1 && CALLS="${CALLS:+$CALLS,}$C"; done
      [ -n "$CALLS" ] && printf '%s\n' "-a always,exit -F arch=$ARCH -S
$CALLS -F auid>=$UID_MIN -F auid!=unset -F key=perm_mod"
   done >> /etc/audit/rules.d/50-perm_mod.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
```

Default Value:

By default, no audit rule is present to record the `chmod`, `fchmod`, and `fchmodat` system calls.

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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	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.		●	●

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.

Audit:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- 'chown' /etc/audit/rules.d/*.rules
```

Verify a rule equivalent to each line below is present:

```
-a always,exit -F arch=b64 -S chown,fchown,lchown,fchownat -F auid>=1000 -F
auid!=unset -F key=perm_mod
-a always,exit -F arch=b32 -S chown,fchown,lchown,fchownat -F auid>=1000 -F
auid!=unset -F key=perm_mod
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- 'chown'
```

Verify the same rules are loaded; `auditctl -l` displays them in the normalized form below:

```
-a always,exit -F arch=b64 -S chown,fchown,lchown,fchownat -F auid>=1000 -F
auid!=-1 -F key=perm_mod
-a always,exit -F arch=b32 -S chown,fchown,lchown,fchownat -F auid>=1000 -F
auid!=-1 -F key=perm_mod
```

Remediation:

1. Add the audit rule(s) for the `chown`, `fchown`, `lchown`, and `fchownat` system calls. The procedure adds a rule for each audit ABI (b64/b32) the running kernel supports and skips any unsupported ABI or system call, so `augenrules --load` does not error. Run:

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

{
   UID_MIN="$(awk '/^[[:space:]]*UID_MIN/{print $2}' /etc/login.defs)"
   for ARCH in b64 b32; do
      CALLS=""; for C in chown fchown lchown fchownat; do ausyscall "$ARCH"
"$C" >/dev/null 2>&1 && CALLS="${CALLS:+$CALLS,}$C"; done
      [ -n "$CALLS" ] && printf '%s\n' "-a always,exit -F arch=$ARCH -S
$CALLS -F auid>=$UID_MIN -F auid!=unset -F key=perm_mod"
   done >> /etc/audit/rules.d/50-perm_mod.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
```

Default Value:

By default, no audit rule is present to record the `chown`, `fchown`, `lchown`, and `fchownat` system calls.

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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	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.		●	●

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.

Audit:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- 'xattr' /etc/audit/rules.d/*.rules
```

Verify a rule equivalent to each line below is present:

```
-a always,exit -F arch=b64 -S
setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr -F
auid>=1000 -F auid!=unset -F key=perm_mod
-a always,exit -F arch=b32 -S
setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr -F
auid>=1000 -F auid!=unset -F key=perm_mod
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- 'xattr'
```

Verify the same rules are loaded; `auditctl -l` displays them in the normalized form below:

```
-a always,exit -F arch=b64 -S
setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr -F
auid>=1000 -F auid!=-1 -F key=perm_mod
-a always,exit -F arch=b32 -S
setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr -F
auid>=1000 -F auid!=-1 -F key=perm_mod
```

Remediation:

1. Add the audit rule(s) for the `setxattr`, `lsetxattr`, `fsetxattr`, `removexattr`, `lremovexattr`, and `fremovexattr` system calls. The procedure adds a rule for each audit ABI (b64/b32) the running kernel supports and skips any unsupported ABI or system call, so `augenrules --load` does not error. Run:

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

{
   UID_MIN="$(awk '/^[[:space:]]*UID_MIN/{print $2}' /etc/login.defs)"
   for ARCH in b64 b32; do
      CALLS=""; for C in setxattr lsetxattr fsetxattr removexattr
lremovexattr fremovexattr; do ausyscall "$ARCH" "$C" >/dev/null 2>&1 &&
CALLS="${CALLS:+$CALLS,}$C"; done
      [ -n "$CALLS" ] && printf '%s\n' "-a always,exit -F arch=$ARCH -S
$CALLS -F auid>=$UID_MIN -F auid!=unset -F key=perm_mod"
   done >> /etc/audit/rules.d/50-perm_mod.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
```

Default Value:

By default, no audit rule is present to record the `setxattr`, `lsetxattr`, `fsetxattr`, `removexattr`, `lremovexattr`, and `fremovexattr` system calls.

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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	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.		●	●

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.

Audit:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- 'mount' /etc/audit/rules.d/*.rules
```

Verify a rule equivalent to each line below is present:

```
-a always,exit -F arch=b64 -S mount -F auid>=1000 -F auid!=unset -k mounts
-a always,exit -F arch=b32 -S mount -F auid>=1000 -F auid!=unset -k mounts
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- 'mount'
```

Verify the same rules are loaded; `auditctl -l` displays them in the normalized form below:

```
-a always,exit -F arch=b64 -S mount -F auid>=1000 -F auid!=-1 -F key=mounts
-a always,exit -F arch=b32 -S mount -F auid>=1000 -F auid!=-1 -F key=mounts
```

Remediation:

1. Add the audit rule(s) for the `mount` system call. The procedure adds a rule for each audit ABI (b64/b32) the running kernel supports and skips any unsupported ABI or system call, so `augenrules --load` does not error. Run:

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

{
   UID_MIN="$(awk '/^[[:space:]]*UID_MIN/{print $2}' /etc/login.defs)"
   for ARCH in b64 b32; do
      CALLS=""; for C in mount; do ausyscall "$ARCH" "$C" >/dev/null 2>&1 &&
CALLS="${CALLS:+$CALLS,}$C"; done
      [ -n "$CALLS" ] && printf '%s\n' "-a always,exit -F arch=$ARCH -S
$CALLS -F auid>=$UID_MIN -F auid!=unset -k mounts"
   done >> /etc/audit/rules.d/50-mounts.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
```

Default Value:

By default, no audit rule is present to record the `mount` system call.

References:

1. GRID: NIX-00000337
2. NIST SP 800-53 :: AU-12 a
3. NIST SP 800-53A :: AU-12.1 (ii)
4. RHEL 8 STIG GROUP ID: V-230425
5. RHEL 8 STIG RULE ID: SV-230425r627750

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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.		●	●

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:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- '(utmp|wtmp|btmp)' /etc/audit/rules.d/*.rules
```

Verify a watch rule covering each target below is present — in either the modern `-F path=`/`-F dir=` form shown (preferred) or the legacy `-w <path> -p <perms>` form; both are a passing state:

```
-a always,exit -S all -F path=/var/run/utmp -F perm=wa -F key=session
-a always,exit -S all -F path=/var/log/wtmp -F perm=wa -F key=session
-a always,exit -S all -F path=/var/log/btmp -F perm=wa -F key=session
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- '(utmp|wtmp|btmp)'
```

Verify the same watch rules are loaded (either form is accepted):

```
-a always,exit -S all -F path=/var/run/utmp -F perm=wa -F key=session
-a always,exit -S all -F path=/var/log/wtmp -F perm=wa -F key=session
-a always,exit -S all -F path=/var/log/btmp -F perm=wa -F key=session
```

Remediation:

1. Add the audit rule(s) that monitor changes to `/var/run/utmp`, `/var/log/wtmp`, and `/var/log/btmp`. The procedure adds each rule only for targets that exist, so `augenrules --load` does not error. Run:

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

{
   [ -e /var/run/utmp ] && printf '%s\n' "-a always,exit -S all -F
path=/var/run/utmp -F perm=wa -F key=session"
   [ -e /var/log/wtmp ] && printf '%s\n' "-a always,exit -S all -F
path=/var/log/wtmp -F perm=wa -F key=session"
   [ -e /var/log/btmp ] && printf '%s\n' "-a always,exit -S all -F
path=/var/log/btmp -F perm=wa -F key=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
```

Default Value:

By default, no audit rule is present to monitor changes to `/var/run/utmp`, `/var/log/wtmp`, and `/var/log/btmp`.

References:

1. GRID: NIX-00000332
2. NIST SP 800-53 Rev. 5: AU-3, AU-3

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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	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.			●

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.

Audit:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- '(lastlog|faillock)' /etc/audit/rules.d/*.rules
```

Verify a watch rule covering each target below is present — in either the modern `-F path=`/`-F dir=` form shown (preferred) or the legacy `-w <path> -p <perms>` form; both are a passing state:

```
-a always,exit -S all -F path=/var/log/lastlog -F perm=wa -F key=logins
-a always,exit -S all -F path=/var/run/faillock -F perm=wa -F key=logins
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- '(lastlog|faillock)'
```

Verify the same watch rules are loaded (either form is accepted):

```
-a always,exit -S all -F path=/var/log/lastlog -F perm=wa -F key=logins
-a always,exit -S all -F path=/var/run/faillock -F perm=wa -F key=logins
```

Remediation:

1. Add the audit rule(s) that monitor changes to `/var/log/lastlog` and `/var/run/faillock`. The procedure adds each rule only for targets that exist, so `augenrules --load` does not error. Run:

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

{
   [ -e /var/log/lastlog ] && printf '%s\n' "-a always,exit -S all -F
path=/var/log/lastlog -F perm=wa -F key=logins"
   [ -e /var/run/faillock ] && printf '%s\n' "-a always,exit -S all -F
path=/var/run/faillock -F perm=wa -F key=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
```

Default Value:

By default, no audit rule is present to monitor changes to `/var/log/lastlog` and `/var/run/faillock`.

References:

1. GRID: NIX-00000329
2. NIST SP 800-53 :: AU-12 a
3. NIST SP 800-53A :: AU-12.1 (ii)
4. RHEL 8 STIG Vul ID: V-230467
5. RHEL 8 STIG Rule ID: SV-230467r627750

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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	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.			●

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.

Audit:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- 'unlink' /etc/audit/rules.d/*.rules
```

Verify a rule equivalent to each line below is present:

```
-a always,exit -F arch=b64 -S unlink,unlinkat -F auid>=1000 -F auid!=unset -F
key=delete
-a always,exit -F arch=b32 -S unlink,unlinkat -F auid>=1000 -F auid!=unset -F
key=delete
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- 'unlink'
```

Verify the same rules are loaded; `auditctl -l` displays them in the normalized form below:

```
-a always,exit -F arch=b64 -S unlink,unlinkat -F auid>=1000 -F auid!=-1 -F
key=delete
-a always,exit -F arch=b32 -S unlink,unlinkat -F auid>=1000 -F auid!=-1 -F
key=delete
```

Remediation:

1. Add the audit rule(s) for the `unlink` and `unlinkat` system calls. The procedure adds a rule for each audit ABI (b64/b32) the running kernel supports and skips any unsupported ABI or system call, so `augenrules --load` does not error. Run:

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

{
   UID_MIN="$(awk '/^[[:space:]]*UID_MIN/{print $2}' /etc/login.defs)"
   for ARCH in b64 b32; do
      CALLS=""; for C in unlink unlinkat; do ausyscall "$ARCH" "$C"
>/dev/null 2>&1 && CALLS="${CALLS:+$CALLS,}$C"; done
      [ -n "$CALLS" ] && printf '%s\n' "-a always,exit -F arch=$ARCH -S
$CALLS -F auid>=$UID_MIN -F auid!=unset -F key=delete"
   done >> /etc/audit/rules.d/50-delete.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
```

Default Value:

By default, no audit rule is present to record the `unlink` and `unlinkat` system calls.

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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.	●	●	●

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.

Audit:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- 'rename' /etc/audit/rules.d/*.rules
```

Verify a rule equivalent to each line below is present:

```
-a always,exit -F arch=b64 -S rename,renameat,renameat2 -F auid>=1000 -F
auid!=unset -F key=delete
-a always,exit -F arch=b32 -S rename,renameat,renameat2 -F auid>=1000 -F
auid!=unset -F key=delete
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- 'rename'
```

Verify the same rules are loaded; `auditctl -l` displays them in the normalized form below:

```
-a always,exit -F arch=b64 -S rename,renameat,renameat2 -F auid>=1000 -F
auid!=-1 -F key=delete
-a always,exit -F arch=b32 -S rename,renameat,renameat2 -F auid>=1000 -F
auid!=-1 -F key=delete
```

Remediation:

1. Add the audit rule(s) for the `rename`, `renameat`, and `renameat2` system calls. The procedure adds a rule for each audit ABI (b64/b32) the running kernel supports and skips any unsupported ABI or system call, so `augenrules --load` does not error. Run:

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

{
   UID_MIN="$(awk '/^[[:space:]]*UID_MIN/{print $2}' /etc/login.defs)"
   for ARCH in b64 b32; do
      CALLS=""; for C in rename renameat renameat2; do ausyscall "$ARCH" "$C"
>/dev/null 2>&1 && CALLS="${CALLS:+$CALLS,}$C"; done
      [ -n "$CALLS" ] && printf '%s\n' "-a always,exit -F arch=$ARCH -S
$CALLS -F auid>=$UID_MIN -F auid!=unset -F key=delete"
   done >> /etc/audit/rules.d/50-delete.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
```

Default Value:

By default, no audit rule is present to record the `rename`, `renameat`, and `renameat2` system calls.

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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.	●	●	●

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 SELinux, an implementation of mandatory access controls. The parameters below monitor any write access (potential additional, deletion or modification of files in the directory) or attribute changes to the `/etc/selinux/`, `/usr/share/selinux/`, and `/var/lib/selinux/` directories.

Note: If a different Mandatory Access Control method is used, changes to the corresponding directories should be audited.

Rationale:

Changes to files in the `/etc/selinux/`, `/usr/share/selinux/`, and `/var/lib/selinux/` directories could indicate that an unauthorized user is attempting to modify access controls and change security contexts, leading to a compromise of the system.

Audit:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- 'selinux' /etc/audit/rules.d/*.rules
```

Verify a watch rule covering each target below is present — in either the modern `-F path=`/`-F dir=` form shown (preferred) or the legacy `-w <path> -p <perms>` form; both are a passing state:

```
-a always,exit -S all -F path=/etc/selinux -F perm=wa -F key=MAC-policy
-a always,exit -S all -F path=/usr/share/selinux -F perm=wa -F key=MAC-policy
-a always,exit -S all -F path=/var/lib/selinux -F perm=wa -F key=MAC-policy
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- 'selinux'
```

Verify the same watch rules are loaded (either form is accepted):

```
-a always,exit -S all -F path=/etc/selinux -F perm=wa -F key=MAC-policy
-a always,exit -S all -F path=/usr/share/selinux -F perm=wa -F key=MAC-policy
-a always,exit -S all -F path=/var/lib/selinux -F perm=wa -F key=MAC-policy
```

Remediation:

1. Add the audit rule(s) that monitor changes to `/etc/selinux`, `/usr/share/selinux`, and `/var/lib/selinux`. The procedure adds each rule only for targets that exist, so `augenrules --load` does not error. Run:

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

{
   [ -e /etc/selinux ] && printf '%s\n' "-a always,exit -S all -F
path=/etc/selinux -F perm=wa -F key=MAC-policy"
   [ -e /usr/share/selinux ] && printf '%s\n' "-a always,exit -S all -F
path=/usr/share/selinux -F perm=wa -F key=MAC-policy"
   [ -e /var/lib/selinux ] && printf '%s\n' "-a always,exit -S all -F
path=/var/lib/selinux -F perm=wa -F key=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
```

Default Value:

By default, no audit rule is present to monitor changes to `/etc/selinux`, `/usr/share/selinux`, and `/var/lib/selinux`.

References:

1. GRID: NIX-00000321
2. NIST SP 800-53 Rev. 5: AU-3, CM-6

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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	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.		●	●

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).

Audit:

1. Run the following command to verify the rule is present in a rules file:

```
# grep -Ps -- 'chcon' /etc/audit/rules.d/*.rules
```

Verify a watch rule covering the command below is present — in either the modern `-F path=` form shown (preferred) or the legacy `-w <path> -p <perms>` form; both are a passing state:

```
-a always,exit -S all -F path=/usr/bin/chcon -F perm=x -F auid>=1000 -F
auid!=unset -k perm_chng
```

2. Run the following command to verify the rule is loaded in the running configuration:

```
# auditctl -l | grep -Ps -- 'chcon'
```

Verify the same watch rule is loaded (either form is accepted; `auditctl -l` shows the kernel-normalized values below):

```
-a always,exit -S all -F path=/usr/bin/chcon -F perm=x -F auid>=1000 -F
auid!=-1 -F key=perm_chng
```

Remediation:

1. Add the audit rule(s) that monitor use of the `/usr/bin/chcon` command. The procedure adds the rule only when the command exists, so `augenrules --load` does not error. Run:

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

{
   UID_MIN="$(awk '/^[[:space:]]*UID_MIN/{print $2}' /etc/login.defs)"
   [ -e /usr/bin/chcon ] && printf '%s\n' "-a always,exit -S all -F
path=/usr/bin/chcon -F perm=x -F auid>=$UID_MIN -F auid!=unset -k perm_chng"
} >> /etc/audit/rules.d/50-perm_chng.rules
```

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
```

Default Value:

By default, no audit rule is present to monitor use of the `/usr/bin/chcon` command.

References:

1. GRID: NIX-00000334
2. NIST SP 800-53 :: AU-12 a
3. NIST SP 800-53A :: AU-12.1 (ii)
4. RHEL 8 STIG Vul ID: V-230419
5. RHEL 8 STIG Rule ID: SV-230419r627750

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) configuration.

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.	●	●	●

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:

1. Run the following command to verify the rule is present in a rules file:

```
# grep -Ps -- 'setfacl' /etc/audit/rules.d/*.rules
```

Verify a watch rule covering the command below is present — in either the modern `-F path=` form shown (preferred) or the legacy `-w <path> -p <perms>` form; both are a passing state:

```
-a always,exit -S all -F path=/usr/bin/setfacl -F perm=x -F auid>=1000 -F
auid!=unset -k perm_chng
```

2. Run the following command to verify the rule is loaded in the running configuration:

```
# auditctl -l | grep -Ps -- 'setfacl'
```

Verify the same watch rule is loaded (either form is accepted; `auditctl -l` shows the kernel-normalized values below):

```
-a always,exit -S all -F path=/usr/bin/setfacl -F perm=x -F auid>=1000 -F
auid!=-1 -F key=perm_chng
```

Remediation:

1. Add the audit rule(s) that monitor use of the `/usr/bin/setfacl` command. The procedure adds the rule only when the command exists, so `augenrules --load` does not error. Run:

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

{
   UID_MIN="$(awk '/^[[:space:]]*UID_MIN/{print $2}' /etc/login.defs)"
   [ -e /usr/bin/setfacl ] && printf '%s\n' "-a always,exit -S all -F
path=/usr/bin/setfacl -F perm=x -F auid>=$UID_MIN -F auid!=unset -k
perm_chng"
} >> /etc/audit/rules.d/50-perm_chng.rules
```

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
```

Default Value:

By default, no audit rule is present to monitor use of the `/usr/bin/setfacl` command.

References:

1. GRID: NIX-00000335
2. NIST SP 800-53 Rev. 5: AU-2, AU-12, SI-5

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) configuration.

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.	●	●	●

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).

Audit:

1. Run the following command to verify the rule is present in a rules file:

```
# grep -Ps -- 'chacl' /etc/audit/rules.d/*.rules
```

Verify a watch rule covering the command below is present — in either the modern `-F path=` form shown (preferred) or the legacy `-w <path> -p <perms>` form; both are a passing state:

```
-a always,exit -S all -F path=/usr/bin/chacl -F perm=x -F auid>=1000 -F
auid!=unset -k perm_chng
```

2. Run the following command to verify the rule is loaded in the running configuration:

```
# auditctl -l | grep -Ps -- 'chacl'
```

Verify the same watch rule is loaded (either form is accepted; `auditctl -l` shows the kernel-normalized values below):

```
-a always,exit -S all -F path=/usr/bin/chacl -F perm=x -F auid>=1000 -F
auid!=-1 -F key=perm_chng
```

Remediation:

1. Add the audit rule(s) that monitor use of the `/usr/bin/chacl` command. The procedure adds the rule only when the command exists, so `augenrules --load` does not error. Run:

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

{
   UID_MIN="$(awk '/^[[:space:]]*UID_MIN/{print $2}' /etc/login.defs)"
   [ -e /usr/bin/chacl ] && printf '%s\n' "-a always,exit -S all -F
path=/usr/bin/chacl -F perm=x -F auid>=$UID_MIN -F auid!=unset -k perm_chng"
} >> /etc/audit/rules.d/50-perm_chng.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
```

Default Value:

By default, no audit rule is present to monitor use of the `/usr/bin/chacl` command.

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. RHEL 8 STIG Vul ID: V-230464
5. RHEL 8 STIG Rule ID: SV-230464r627750

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) configuration.

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.	●	●	●

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).

Audit:

1. Run the following command to verify the rule is present in a rules file:

```
# grep -Ps -- 'usermod' /etc/audit/rules.d/*.rules
```

Verify a watch rule covering the command below is present — in either the modern `-F path=` form shown (preferred) or the legacy `-w <path> -p <perms>` form; both are a passing state:

```
-a always,exit -S all -F path=/usr/sbin/usermod -F perm=x -F auid>=1000 -F
auid!=unset -k usermod
```

2. Run the following command to verify the rule is loaded in the running configuration:

```
# auditctl -l | grep -Ps -- 'usermod'
```

Verify the same watch rule is loaded (either form is accepted; `auditctl -l` shows the kernel-normalized values below):

```
-a always,exit -S all -F path=/usr/sbin/usermod -F perm=x -F auid>=1000 -F
auid!=-1 -F key=usermod
```

Remediation:

1. Add the audit rule(s) that monitor use of the `/usr/sbin/usermod` command. The procedure adds the rule only when the command exists, so `augenrules --load` does not error. Run:

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

{
   UID_MIN="$(awk '/^[[:space:]]*UID_MIN/{print $2}' /etc/login.defs)"
   [ -e /usr/sbin/usermod ] && printf '%s\n' "-a always,exit -S all -F
path=/usr/sbin/usermod -F perm=x -F auid>=$UID_MIN -F auid!=unset -k usermod"
} >> /etc/audit/rules.d/50-usermod.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
```

Default Value:

By default, no audit rule is present to monitor use of the `/usr/sbin/usermod` command.

References:

1. GRID: NIX-00000336
2. NIST SP 800-53 :: AU-12 a
3. NIST SP 800-53A :: AU-12.1 (ii)
4. RHEL 8 STIG Vul ID: V-230463
5. RHEL 8 STIG Rule ID: SV-230463r627750

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) configuration.

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.	●	●	●

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.

Audit:

1. Run the following command to verify the rule is present in a rules file:

```
# grep -Ps -- '/usr/bin/kmod' /etc/audit/rules.d/*.rules
```

Verify a watch rule covering the command below is present — in either the modern `-F path=` form shown (preferred) or the legacy `-w <path> -p <perms>` form; both are a passing state:

```
-a always,exit -S all -F path=/usr/bin/kmod -F perm=x -F auid>=1000 -F
auid!=unset -k kernel_modules
```

2. Run the following command to verify the rule is loaded in the running configuration:

```
# auditctl -l | grep -Ps -- '/usr/bin/kmod'
```

Verify the same watch rule is loaded (either form is accepted; `auditctl -l` shows the kernel-normalized values below):

```
-a always,exit -S all -F path=/usr/bin/kmod -F perm=x -F auid>=1000 -F
auid!=-1 -F key=kernel_modules
```

Remediation:

1. Add the audit rule(s) that monitor use of the `/usr/bin/kmod` command. The procedure adds each rule only for targets that exist, so `augenrules --load` does not error. Run:

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

{
   UID_MIN="$(awk '/^[[:space:]]*UID_MIN/{print $2}' /etc/login.defs)"
   [ -e /usr/bin/kmod ] && printf '%s\n' "-a always,exit -S all -F
path=/usr/bin/kmod -F perm=x -F auid>=$UID_MIN -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
```

Default Value:

By default, no audit rule is present to monitor use of the `/usr/bin/kmod` command.

References:

1. GRID: NIX-00000327
2. NIST SP 800-53 :: AU-12 a
3. NIST SP 800-53A :: AU-12.1 (ii)
4. RHEL 8 STIG Vul ID: V-230438
5. RHEL 8 STIG Rule ID: SV-230438r810464
6. RHEL 8 STIG Vul ID: V-230446
7. RHEL 8 STIG Rule ID: SV-230446r627750
8. RHEL 8 STIG Vul ID: V-230465
9. RHEL 8 STIG Rule ID: SV-230465r627750

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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.	●	●	●

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.

Audit:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- 'init_module' /etc/audit/rules.d/*.rules
```

Verify a rule equivalent to each line below is present:

```
-a always,exit -F arch=b64 -S init_module,finit_module -F auid>=1000 -F
auid!=unset -k kernel_modules
-a always,exit -F arch=b32 -S init_module,finit_module -F auid>=1000 -F
auid!=unset -k kernel_modules
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- 'init_module'
```

Verify the same rules are loaded; `auditctl -l` displays them in the normalized form below:

```
-a always,exit -F arch=b64 -S init_module,finit_module -F auid>=1000 -F
auid!=-1 -F key=kernel_modules
-a always,exit -F arch=b32 -S init_module,finit_module -F auid>=1000 -F
auid!=-1 -F key=kernel_modules
```

Remediation:

1. Add the audit rule(s) for the `init_module` and `finit_module` system calls. The procedure adds a rule for each audit ABI (b64/b32) the running kernel supports and skips any unsupported ABI or system call, so `augenrules --load` does not error. Run:

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

{
   UID_MIN="$(awk '/^[[:space:]]*UID_MIN/{print $2}' /etc/login.defs)"
   for ARCH in b64 b32; do
      CALLS=""; for C in init_module finit_module; do ausyscall "$ARCH" "$C"
>/dev/null 2>&1 && CALLS="${CALLS:+$CALLS,}$C"; done
      [ -n "$CALLS" ] && printf '%s\n' "-a always,exit -F arch=$ARCH -S
$CALLS -F auid>=$UID_MIN -F auid!=unset -k kernel_modules"
   done >> /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
```

Default Value:

By default, no audit rule is present to record the `init_module` and `finit_module` system calls.

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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.	●	●	●

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.

Audit:

1. Run the following command to verify the rule(s) are present in a rules file:

```
# grep -Ps -- 'delete_module' /etc/audit/rules.d/*.rules
```

Verify a rule equivalent to each line below is present:

```
-a always,exit -F arch=b64 -S delete_module -F auid>=1000 -F auid!=unset -k
kernel_modules
-a always,exit -F arch=b32 -S delete_module -F auid>=1000 -F auid!=unset -k
kernel_modules
```

2. Run the following command to verify the rule(s) are loaded in the running configuration:

```
# auditctl -l | grep -Ps -- 'delete_module'
```

Verify the same rules are loaded; `auditctl -l` displays them in the normalized form below:

```
-a always,exit -F arch=b64 -S delete_module -F auid>=1000 -F auid!=-1 -F
key=kernel_modules
-a always,exit -F arch=b32 -S delete_module -F auid>=1000 -F auid!=-1 -F
key=kernel_modules
```

Remediation:

1. Add the audit rule(s) for the `delete_module` system call. The procedure adds a rule for each audit ABI (b64/b32) the running kernel supports and skips any unsupported ABI or system call, so `augenrules --load` does not error. Run:

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

{
   UID_MIN="$(awk '/^[[:space:]]*UID_MIN/{print $2}' /etc/login.defs)"
   for ARCH in b64 b32; do
      CALLS=""; for C in delete_module; do ausyscall "$ARCH" "$C" >/dev/null
2>&1 && CALLS="${CALLS:+$CALLS,}$C"; done
      [ -n "$CALLS" ] && printf '%s\n' "-a always,exit -F arch=$ARCH -S
$CALLS -F auid>=$UID_MIN -F auid!=unset -k kernel_modules"
   done >> /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
```

Default Value:

By default, no audit rule is present to record the `delete_module` system call.

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) 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.	●	●	●

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
```

Default Value:

By default, the `-c` directive is not configured; audit rule loading halts when an error is encountered.

References:

1. GRID: NIX-00000341

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.		●	●

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.

Impact:

Once the audit configuration is made immutable, any change to audit rules requires a system reboot to take effect.

Audit:

1. Run the following command to verify the immutable flag is set in a rules file. This on-disk setting is the persistent compliance determinant:

```
# grep -Ph -- '^\h*-e\h+2\b' /etc/audit/rules.d/*.rules | tail -1

-e 2
```

2. Run the following command to check whether the immutable flag is active in the running configuration:

```
# auditctl -s | grep 'enabled'

enabled 2
```

A running value of `enabled 2` confirms the configuration is currently immutable. If step 1 returns `-e 2` but this returns `enabled 1`, the immutable flag is configured on disk but not yet active in the running kernel; it takes effect when the audit rules are reloaded at the next reboot. This is expected immediately after applying the audit rules (immutability is finalized on reboot so that rule changes are not locked out mid-configuration) and is not, on its own, a finding.

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 '%s\n' "" "-e 2" >> /etc/audit/rules.d/99-finalize.rules
```

Load audit rules

Merge and load the rules into active configuration:

```
# augenrules --load
```

Check if reboot is required.

```
# if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then printf "Reboot
required to load rules\n"; fi
```

Default Value:

By default, the audit configuration is not immutable (the `-e 2` directive is not set).

References:

1. GRID: NIX-00000340
2. NIST SP 800-53 :: AU-9
3. NIST SP 800-53A :: AU-9.1
4. RHEL 8 STIG Vul ID: V-230402
5. RHEL 8 STIG Rule ID: SV-230402r627750

Additional Information:

Refer to the "Configure auditd Rules" section overview for general guidance applicable to all recommendations in this section: rule syntax, the `arch=` field, the deprecated `-w` form, `auid>=`/UID_MIN, watch-target existence, system-call line structure, and reboot behavior with an immutable (`-e 2`) configuration.

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.		●	●

6.2.3.36 Ensure the running and on disk configuration is the same (Automated)

Profile Applicability:

- Level 2 - Server
- Level 2 - Workstation

Description:

The Audit system have both on disk and running configuration. It is possible for these configuration settings to differ.

Note: Due to the limitations of `augenrules` and `auditctl`, it is not absolutely guaranteed that loading the rule sets via `augenrules --load` will result in all rules being loaded or even that the user will be informed if there was a problem loading the rules.

Rationale:

Configuration differences between what is currently running and what is on disk could cause unexpected problems or may give a false impression of compliance requirements.

Audit:

Merged rule sets

Ensure that all rules in `/etc/audit/rules.d` have been merged into `/etc/audit/audit.rules`:

```
# augenrules --check

/usr/sbin/augenrules: No change
```

Should there be any drift, run `augenrules --load` to merge and load all rules.

Remediation:

If the rules are not aligned across all three areas (the `*.rules` files in `/etc/audit/rules.d/`, the merged `/etc/audit/audit.rules`, and the running kernel configuration), run the following command to merge and load all rules:

```
# augenrules --load
```

Check if reboot is required.

```
if [[ $(auditctl -s | grep "enabled") =~ "2" ]]; then echo "Reboot required
to load rules"; fi
```

References:

1. 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.		●	●

6.2.4 Configure auditd File Access

This subsection restricts access to the components of the audit system: the audit **log files** and their directory (derived from `log_file` in `/etc/audit/auditd.conf`), the audit **configuration files** (`*.conf` and `*.rules` under `/etc/audit/`), and the audit **tools** (`auditctl`, `aureport`, `ausearch`, `autrace`, `auditd`, `augenrules`). Each is checked for an appropriately restrictive mode and correct user/group ownership so that unauthorized users cannot read, alter, or remove audit data, tamper with audit configuration, or misuse the audit utilities. The audit log group is typically `adm` (or `root`).

6.2.4.1 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.

Impact:

Overly restrictive permissions on the audit log directory could prevent processes or the log-management group from accessing the audit logs.

Audit:

Run the following command to verify the audit log file directory is mode `0750` or more restrictive:

```
# stat -Lc '%n %a' "$(dirname "$(awk -F= '/^\s*log_file\s*/{print $2}'
/etc/audit/auditd.conf | xargs)")"
```

Example output:

```
/var/log/audit 750
```

The mode should be `0750` or more restrictive (no write access for group and no access for other).

Remediation:

Run the following command to configure the audit log directory to have a mode of "0750" or more restrictive:

```
# chmod g-w,o-rwx "$(dirname "$(awk -F= '/^\s*log_file\s*/{print $2}'
/etc/audit/auditd.conf | xargs)")"
```

Default Value:

`0700`

References:

1. GRID: NIX-00000342
2. NIST SP 800-53 :: AU-9
3. NIST SP 800-53A :: AU-9.1 (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.	●	●	●
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.	●	●	●

6.2.4.2 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.

Impact:

Restricting audit log file permissions to the recommended mode may prevent tools or users that previously relied on broader access from reading the logs.

Audit:

Run the following command to verify all audit log files are mode `0640` or more restrictive:

```
# [ -f /etc/audit/auditd.conf ] && find "$(dirname "$(awk -F=
'/^\s*log_file\s*/{print $2}' /etc/audit/auditd.conf | xargs)")" -maxdepth 1
-type f -perm /0137
```

Nothing should be returned.

Remediation:

Run the following command to remove more permissive mode than `0640` from audit log files:

```
# [ -f /etc/audit/auditd.conf ] && find "$(dirname $(awk -F "="
'/^\s*log_file/ {print $2}' /etc/audit/auditd.conf | xargs))" -type f -perm
/0137 -exec chmod u-x,g-wx,o-rwx {} +
```

Default Value:

`0600`

References:

1. GRID: NIX-00000291
2. NIST SP 800-53 :: AU-9
3. NIST SP 800-53A :: AU-9.1 (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.	●	●	●
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.	●	●	●

6.2.4.3 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.

Impact:

Restricting audit log file ownership to `root` may prevent non-root accounts that previously read the logs from accessing them.

Audit:

Run the following command to verify all audit log files are owned by the `root` user:

```
# [ -f /etc/audit/auditd.conf ] && find "$(dirname "$(awk -F=
'/^\s*log_file\s*/{print $2}' /etc/audit/auditd.conf | xargs)")" -maxdepth 1
-type f ! -user root
```

Nothing should be returned.

Remediation:

Run the following command to configure the audit log files to be owned by the `root` user:

```
# [ -f /etc/audit/auditd.conf ] && find "$(dirname $(awk -F "="
'/^\s*log_file/ {print $2}' /etc/audit/auditd.conf | xargs))" -type f ! -user
root -exec chown root {} +
```

Default Value:

`root`

References:

1. GRID: NIX-00000292
2. NIST SP 800-53 :: AU-9
3. NIST SP 800-53A :: AU-9.1 (ii)
4. RHEL 8 STIG Vul ID: V-230397
5. RHEL 8 STIG Rule ID: SV-230397r627750

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.	●	●	●

6.2.4.4 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.

Impact:

Restricting audit log file group ownership to `adm` (or `root`) may prevent groups that previously relied on access from reading the logs.

Audit:

Both of the following must be true.

1. Run the following command to verify the `log_group` parameter is set to `root` or `adm`:

```
# grep -Po '^\s*log_group\s*=\s*\K\S+' /etc/audit/auditd.conf
```

The returned value should be `root` or `adm`.

2. Run the following command to verify all audit log files are group-owned by `root` or `adm`:

```
# [ -f /etc/audit/auditd.conf ] && find "$(dirname "$(awk -F=
'/^\s*log_file\s*/{print $2}' /etc/audit/auditd.conf | xargs)")" -maxdepth 1
-type f \( ! -group root -a ! -group adm \)
```

Nothing should be returned.

Remediation:

1. Run the following command to configure the audit log files to be owned by `adm` group:

```
# find $(dirname $(awk -F"=" '/^\s*log_file\s*=\s*/ {print $2}'
/etc/audit/auditd.conf | xargs)) -type f \( ! -group adm -a ! -group root \)
-exec chgrp adm {} +
```

2. Run the following command to set the `log_group` parameter in the audit configuration file to `log_group = adm`:

```
# sed -ri 's/^\s*#?\s*log_group\s*=\s*\S+(\s*#.*)?.*$/log_group = adm\1/'
/etc/audit/auditd.conf
```

3. Run the following command to restart the audit daemon to reload the configuration file:

```
# service auditd restart
```

Note: If the audit log file directory is not configured (`log_file` is unset in `/etc/audit/auditd.conf`), set a valid `log_file` location before applying the steps above.

Default Value:

`root`

References:

1. GRID: NIX-00000290
2. NIST SP 800-53 :: AU-9
3. NIST SP 800-53A :: AU-9.1 (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.	●	●	●
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.	●	●	●

6.2.4.5 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.

Impact:

Restricting audit configuration file permissions to the recommended mode may prevent tools or users that previously relied on broader access.

Audit:

Run the following command to verify all audit configuration files are mode `0640` or more restrictive:

```
# find /etc/audit/ -type f \( -name '*.conf' -o -name '*.rules' \) -perm
/0137
```

Nothing should be returned.

Remediation:

Run the following command to remove more permissive mode than 0640 from the audit configuration files:

```
# find /etc/audit/ -type f \( -name '*.conf' -o -name '*.rules' \) -exec
chmod u-x,g-wx,o-rwx {} +
```

Default Value:

`0640`

References:

1. GRID: NIX-00000288
2. NIST SP 800-53 :: AU-9
3. NIST SP 800-53A :: AU-9.1 (ii)
4. RHEL 8 STIG Vul ID: V-230471
5. RHEL 8 STIG Rule ID: SV-230471r627750

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.	●	●	●

6.2.4.6 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.

Impact:

Restricting audit configuration file ownership to `root` may prevent non-root accounts that previously accessed them.

Audit:

Run the following command to verify all 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 {} +
```

Default Value:

`root`

References:

1. GRID: NIX-00000289
2. NIST SP 800-53 :: AU-9
3. NIST SP 800-53A :: AU-9.1 (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.	●	●	●
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.	●	●	●

6.2.4.7 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.

Impact:

Restricting audit configuration file group ownership to `root` may prevent groups that previously accessed them.

Audit:

Run the following command to verify all audit configuration files are group-owned by the `root` group:

```
# find /etc/audit/ -type f \( -name '*.conf' -o -name '*.rules' \) ! -group
root
```

Nothing should be returned.

Remediation:

Run the following command to change group to `root`:

```
# find /etc/audit/ -type f \( -name '*.conf' -o -name '*.rules' \) ! -group
root -exec chgrp root {} +
```

Default Value:

`root`

References:

1. GRID: NIX-00000287
2. NIST SP 800-53 :: AU-9
3. NIST SP 800-53A :: AU-9.1 (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.	●	●	●
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.	●	●	●

6.2.4.8 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.

Impact:

Removing group and other write access to the audit tools prevents unauthorized modification; normal use of the tools is unaffected.

Audit:

Run the following command to verify all audit tools are mode `0755` or more restrictive:

```
# stat -Lc '%n %a' /sbin/auditctl /sbin/aureport /sbin/ausearch /sbin/autrace
/sbin/auditd /sbin/augenrules
```

Example output:

```
/sbin/auditctl 755
/sbin/aureport 755
/sbin/ausearch 755
/sbin/autrace 755
/sbin/auditd 755
/sbin/augenrules 755
```

Each audit tool should be mode `0755` or more restrictive (no write access for group or other).

Remediation:

Run the following command to remove more permissive mode from the audit tools:

```
# chmod go-w /sbin/auditctl /sbin/aureport /sbin/ausearch /sbin/autrace
/sbin/auditd /sbin/augenrules
```

Default Value:

`0755`

References:

1. GRID: NIX-00000297
2. NIST SP 800-53 :: AU-9
3. NIST SP 800-53A :: AU-9.1 (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.	●	●	●
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.	●	●	●

6.2.4.9 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.

Impact:

Setting audit tool ownership to `root` prevents non-root accounts from modifying them; normal use is unaffected.

Audit:

Run the following command to verify all audit tools are owned by the `root` user:

```
# stat -Lc "%n %U" /sbin/auditctl /sbin/aureport /sbin/ausearch /sbin/autrace
/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/autrace
/sbin/auditd /sbin/augenrules
```

Default Value:

`root`

References:

1. GRID: NIX-00000298
2. NIST SP 800-53 :: AU-9
3. NIST SP 800-53A :: AU-9.1 (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.	●	●	●
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.	●	●	●

6.2.4.10 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.

Impact:

Setting audit tool group ownership to `root` prevents unauthorized group modification; normal use is unaffected.

Audit:

Run the following command to verify all audit tools are group-owned by the `root` group:

```
# stat -Lc "%n %G" /sbin/auditctl /sbin/aureport /sbin/ausearch /sbin/autrace
/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/autrace
/sbin/auditd /sbin/augenrules
```

Default Value:

`root`

References:

1. GRID: NIX-00000296
2. NIST SP 800-53 :: AU-9
3. NIST SP 800-53A :: AU-9.1 (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.	●	●	●
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.	●	●	●

6.3 Configure Integrity Checking

This subsection configures file integrity monitoring with AIDE (Advanced Intrusion Detection Environment): installing AIDE, scheduling regular integrity checks, and using AIDE to protect the integrity of the audit tools. AIDE cannot prevent intrusions, but it detects unauthorized changes to files by comparing the current filesystem state against a baseline database. When setting up AIDE, decide internally what the site policy will be concerning integrity checking; review the AIDE documentation before proceeding.

Note: If another local site approved solution is used, and configured to provide the capabilities covered by AIDE in this section, this section may be skipped.

6.3.1 Ensure AIDE is installed (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

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

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

Rationale:

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

Impact:

Installing AIDE and initializing its database consumes CPU, I/O, and disk space; the initial `aide --init` may take time on large filesystems.

Audit:

Run the following command to verify the `aide` package is installed:

```
# rpm -q aide

aide-<version>
```

Remediation:

1. Run the following command to install `aide`:

```
# dnf install aide
```

2. Configure `aide` as appropriate for your environment. Consult the `aide` documentation for options.
3. Run the following command to Initialize `aide`:

```
# aide --init
```

4. Run the following command to create a baseline copy of the AIDE database:

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

References:

1. GRID: NIX-00000347
2. NIST SP 800-53 :: SI-7
3. NIST SP 800-53A :: SI-7.1 (ii)
4. RHEL 8 STIG Vul ID: V-251710
5. RHEL 8 STIG Rule ID: SV-251710r880730
6. AIDE stable manual: http://aide.sourceforge.net/stable/manual.html

Additional Information:

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

CIS Controls:

Controls Version	Control	IG 1	IG 2	IG 3
v8	3.14 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).			●

6.3.2 Ensure filesystem integrity is regularly checked (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Periodic checking of filesystem integrity detects unauthorized changes to files. AIDE (Advanced Intrusion Detection Environment) compares the current state of the filesystem against a previously initialized baseline database; this recommendation ensures that comparison runs automatically on a regular schedule rather than only on demand.

Rationale:

Periodic file checking allows the system administrator to determine on a regular basis if critical files have been changed in an unauthorized fashion. Running the integrity check on a defined schedule -- rather than manually or ad hoc -- ensures changes are detected consistently and reduces the window in which unauthorized modifications can go unnoticed.

Impact:

Scheduled integrity checks consume CPU and I/O when they run; choose a schedule and time that balance timely detection against system load.

Audit:

A periodic AIDE integrity check should be scheduled through cron or an enabled or active systemd timer. At least one of the following should return a matching entry.

1. Run the following command to list cron-scheduled AIDE checks:

```
# grep -rEi 'aide[[:space:]]+(--check|--update)' /var/spool/cron/
/etc/crontab /etc/cron.d /etc/cron.hourly /etc/cron.daily /etc/cron.weekly
/etc/cron.monthly 2>/dev/null
```

Example output:

```
/etc/cron.d/aidecheck:0 5 * * * root /usr/sbin/aide --check
```

2. Run the following command to list enabled or active systemd timers that run AIDE:

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

{
   for t in $(systemctl list-unit-files --type=timer --no-legend | awk
'{print $1}'); do
      en="$(systemctl is-enabled "$t" 2>/dev/null)"; ac="$(systemctl is-
active "$t" 2>/dev/null)"
      [ "$en" = "enabled" ] || [ "$ac" = "active" ] || continue
      s="$(systemctl show -p Unit --value "$t")"; [ -z "$s" ] &&
s="${t%.timer}.service"
      systemctl cat "$s" 2>/dev/null | grep -Eq
'ExecStart=.*aide[[:space:]]+(--check|--update)' && echo "$t -> $s"
   done
}
```

Example output:

```
aidecheck.timer -> aidecheck.service
```

At least one of the two commands above should return a matching entry, confirming that a periodic AIDE integrity check is scheduled (via cron or a systemd timer).

Remediation:

Configure a systemd timer to run a periodic AIDE integrity check.

1. Create the file `/etc/systemd/system/aidecheck.service` with the following content:

```
[Unit]
Description=Aide Check

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

2. Create the file `/etc/systemd/system/aidecheck.timer` with the following content:

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

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

[Install]
WantedBy=timers.target
```

3. Run the following commands to set access on the unit files, reload systemd, and enable the timer:

```
# chown root:root /etc/systemd/system/aidecheck.service
/etc/systemd/system/aidecheck.timer
# chmod 0644 /etc/systemd/system/aidecheck.service
/etc/systemd/system/aidecheck.timer
# systemctl daemon-reload
# systemctl enable --now aidecheck.timer
```

References:

1. GRID: NIX-00000349
2. https://github.com/konstruktoid/hardening/blob/master/config/aidecheck.service
3. https://github.com/konstruktoid/hardening/blob/master/config/aidecheck.timer
4. NIST SP 800-53 :: SI-7
5. NIST SP 800-53A :: SI-7.1 (ii)

Additional Information:

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

CIS Controls:

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

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

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Audit tools include the vendor-provided and open-source utilities used to view and manipulate audit records and settings (for example `auditctl`, `auditd`, `ausearch`, `aureport`, `autrace`, and `augenrules`, along with custom queries and report generators). This recommendation uses AIDE (Advanced Intrusion Detection Environment) to protect the integrity of those tools by recording cryptographic hashes and file attributes, so that modification or replacement of a tool can be detected.

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.

Impact:

After a legitimate update to an audit tool (for example, a package update), AIDE will report the change until its database is refreshed (`aide --update`, or re-initialize the database); refresh the AIDE database after patching so the new binaries are accepted as the baseline.

Audit:

Run the following command to verify that AIDE is configured to monitor the integrity of the audit tools:

```
# grep -E '/(auditctl|auditd|ausearch|aureport|autrace|augenrules)\b'
/etc/aide.conf
```

Each of the six audit tools -- `auditctl`, `auditd`, `ausearch`, `aureport`, `autrace`, `augenrules` (at their absolute path, typically under `/usr/sbin`) -- should have a selection line applying at least the attributes `p+i+n+u+g+s+b+acl+xattrs+sha512` (directly or via an equivalent rule group).

Example output:

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

Note: If `/etc/aide.conf` uses an `@@x_include` statement, also check the included configuration files.

Remediation:

1. 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.

2. Edit `/etc/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>/autrace p+i+n+u+g+s+b+acl+xattrs+sha512
<PATH>/augenrules p+i+n+u+g+s+b+acl+xattrs+sha512
```

Example

```
# printf '\n%s' "# 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" "$(readlink -f /sbin/autrace)
p+i+n+u+g+s+b+acl+xattrs+sha512" "$(readlink -f /sbin/augenrules)
p+i+n+u+g+s+b+acl+xattrs+sha512" >> /etc/aide.conf
```

Note: - IF - `/etc/aide.conf` includes a `@@x_include` statement:

Example:

```
@@x_include /etc/aide.conf.d ^[a-zA-Z0-9_-]+$
```

- `@@x_include` FILE
- `@@x_include` DIRECTORY REGEX
    - `@x_include` is identical to `@@include`, except that if a config file is executable it 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 and must not be group or world-writable.

References:

1. GRID: NIX-00000348
2. AIDE.CONF(5)
3. NIST SP 800-53 :: SI-7
4. NIST SP 800-53A :: SI-7.1 (ii)

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).			●

6.3.4 Ensure report_url is configured (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

The report_url directive in aide.conf names the destination(s) to which AIDE writes each integrity check's report. It accepts multiple destinations and supports a syslog:<facility> target - where <facility> is a syslog facility, such as one of LOG_LOCAL0 through LOG_LOCAL7 - that sends the report to the system log rather than, or in addition to, a local file. AIDE has no built-in mail transport; syslog is its native mechanism for directing the report to the system logging service.

Rationale:

A scheduled integrity check is only useful if its findings reach a monitored location. When the report is written only to a local file - or to standard output that the scheduling job discards - an unauthorized addition, modification, or deletion of a critical system file or audit tool can be recorded where no one is watching, and on the same host an attacker may be able to alter or delete it. Writing the report to the system log routes integrity findings into the centrally collected, retained, and access-controlled log stream, so discrepancies are preserved off the checked file's local context and are available for review and alerting.

Impact:

Integrity-check reports written to the system log add to log volume, and a ruleset that has not been tuned to exclude files expected to change frequently can generate sizable or noisy log entries. The system logging service must be configured (see the logging recommendations) for the report to be collected and retained.

Audit:

Run the following command to verify that AIDE writes its integrity-check report to the system log. Review the effective aide.conf, including any files pulled in by `@@include` or `@@x_include` directives:

```
# grep -P '^\h*report_url' /etc/aide.conf
```

A report_url directive sending the report to syslog should be returned, for example:

Example output:

```
report_url=syslog:LOG_LOCAL0
```

Additional report_url destinations (such as report_url=file:/var/log/aide.log) may also be present; at least one syslog: destination should be configured. If nothing is returned, or report_url resolves only to a non-logging destination, AIDE is not sending its report to the system log.

Remediation:

Edit /etc/aide.conf and add a `report_url` directive that writes the check report to the system log, using a syslog facility that the system logging service is configured to collect (the local-use facilities `LOG_LOCAL0` through `LOG_LOCAL7` are commonly used for this):

```
report_url=syslog:<facility>
```

Example:

```
report_url=syslog:LOG_LOCAL0
```

AIDE supports multiple report_url destinations, so an existing destination (such as `report_url=file:/var/log/aide.log`) may be kept alongside the syslog entry. Ensure the system logging service is configured to collect and retain the selected facility (see the logging recommendations) so the report is captured in the central log stream.

References:

1. GRID: NIX-00000430
2. NIST SP 800-53 :: SI-7(2)
3. NIST SP 800-53 :: SI-7
4. aide.conf(5)

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).			●

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.

7.1 Configure system file and directory access

This section provides recommendations for configuring access to the critical account-database and system files: the user/group databases (`/etc/passwd`, `/etc/group`, `/etc/shadow`, `/etc/gshadow`) and their backups, the valid-login-shells list (`/etc/shells`), and the password-history store (`/etc/security/opasswd`); it also covers locating world-writable, unowned/ungrouped, and SUID/SGID files. The backup copies (the `-`-suffixed files and `opasswd.old`) are written by account-management tools and may not be present on every system; where such a file does not exist, the corresponding recommendation is not applicable and passes.

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 that access to `/etc/passwd` is configured (mode `0644` or more restrictive, owned by `root:root`):

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

The following should be returned:

```
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 :: 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.	●	●	●

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 that access to `/etc/passwd-` is configured. `/etc/passwd-` is a backup written by account-management tools and may not exist:

```
# [ -e /etc/passwd- ] && stat -Lc 'Access: (%#a/%A)  Uid: ( %u/ %U) Gid: (
%g/ %G)' /etc/passwd-
```

If the file exists, the following should be returned; **- OR -** if it does not exist, nothing is returned (the recommendation is not applicable):

```
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 :: 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.	●	●	●

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.

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 that access to `/etc/group` is configured (mode `0644` or more restrictive, owned by `root:root`):

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

The following should be returned:

```
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 :: 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.	●	●	●

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 that access to `/etc/group-` is configured. `/etc/group-` is a backup written by account-management tools and may not exist:

```
# [ -e /etc/group- ] && stat -Lc 'Access: (%#a/%A)  Uid: ( %u/ %U) Gid: ( %g/
%G)' /etc/group-
```

If the file exists, the following should be returned; **- OR -** if it does not exist, nothing is returned (the recommendation is not applicable):

```
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 :: 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.	●	●	●

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 that access to `/etc/shadow` is configured (mode `0000`, owned by `root:root`):

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

The following should be returned:

```
Access: (0000/----------)  Uid: (0/root)   Gid: (0/root)
```

Remediation:

Run the following commands to set mode, owner, and group on `/etc/shadow`:

```
# chown root:root /etc/shadow
# chmod 0000 /etc/shadow
```

Default Value:

Access: (0/----------) Uid: (0/ root) Gid: (0/ root)

References:

1. GRID: NIX-00000357
2. NIST SP 800-53 :: 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.	●	●	●

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 that access to `/etc/shadow-` is configured. `/etc/shadow-` is a backup written by account-management tools and may not exist:

```
# [ -e /etc/shadow- ] && stat -Lc 'Access: (%#a/%A)  Uid: ( %u/ %U) Gid: (
%g/ %G)' /etc/shadow-
```

If the file exists, the following should be returned; **- OR -** if it does not exist, nothing is returned (the recommendation is not applicable):

```
Access: (0000/----------)  Uid: (0/root)   Gid: (0/root)
```

Remediation:

Run the following commands to set mode, owner, and group on `/etc/shadow-`:

```
# chown root:root /etc/shadow-
# chmod 0000 /etc/shadow-
```

Default Value:

Access: (0/----------) Uid: (0/ root) Gid: (0/ root)

References:

1. GRID: NIX-00000358
2. NIST SP 800-53 :: 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.	●	●	●

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 that access to `/etc/gshadow` is configured (mode `0000`, owned by `root:root`):

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

The following should be returned:

```
Access: (0000/----------)  Uid: (0/root)   Gid: (0/root)
```

Remediation:

Run the following commands to set mode, owner, and group on `/etc/gshadow`:

```
# chown root:root /etc/gshadow
# chmod 0000 /etc/gshadow
```

Default Value:

Access: (0/----------) Uid: (0/ root) Gid: (0/ root)

References:

1. GRID: NIX-00000352
2. NIST SP 800-53 :: 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.	●	●	●

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 that access to `/etc/gshadow-` is configured. `/etc/gshadow-` is a backup written by account-management tools and may not exist:

```
# [ -e /etc/gshadow- ] && stat -Lc 'Access: (%#a/%A)  Uid: ( %u/ %U) Gid: (
%g/ %G)' /etc/gshadow-
```

If the file exists, the following should be returned; **- OR -** if it does not exist, nothing is returned (the recommendation is not applicable):

```
Access: (0000/----------)  Uid: (0/root)   Gid: (0/root)
```

Remediation:

Run the following commands to set mode, owner, and group on `/etc/gshadow-`:

```
# chown root:root /etc/gshadow-
# chmod 0000 /etc/gshadow-
```

Default Value:

Access: (0/----------) Uid: (0/ root) Gid: (0/ root)

References:

1. GRID: NIX-00000353
2. NIST SP 800-53 :: 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.	●	●	●

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 that access to `/etc/shells` is configured (mode `0644` or more restrictive, owned by `root:root`):

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

The following should be returned:

```
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 :: 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.	●	●	●

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:

```
# [ -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 :: 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.	●	●	●

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.

Impact:

Removing write access from world-writable files, or adding the sticky bit to world-writable directories, can break applications that depend on that access; consult vendor documentation before applying.

Audit:

Run the following script to list any world-writable files and any world-writable directories that do not have the sticky bit set:

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

{
   a_path=(-path "*/containers/storage/*" -o -path "*/containerd/*" -o -path
"*/kubelet/*" -o -path "*/docker/*" -o -path "/sys/*" -o -path "/snap/*" -o -
path "/boot/efi/*")
   while IFS= read -r l_mount; do
      find "$l_mount" -mount -xdev \( "${a_path[@]}" \) -prune -o \
         \( \( -type f -perm -0002 \) -o \( -type d -perm -0002 ! -perm -1000
\) \) \
         -printf '%y %p\n' 2>/dev/null
   done < <(findmnt -Dkerno fstype,target | awk '($1 !~
/^[[:space:]]*(nfs|proc|cifs|smb|vfat|iso9660|efivarfs|selinuxfs|ncpfs)/ &&
$2 !~ /^(\/run|\/tmp|\/var\/tmp)(\/|$)/){print $2}')
}
```

Nothing should be returned (each result is tagged `f` = file or `d` = directory).

Note: on systems with many files/directories this may run for some time.

Remediation:

- World Writable Files:
  - It is recommended that write access is removed from `other` with the command (`chmod o-w <filename>`), but always consult relevant vendor documentation to avoid breaking any application dependencies on a given file.
- World Writable Directories:
  - Set the sticky bit on all world writable directories with the command (`chmod a+t <directory_name>`)

Run the following script to:

- Remove other write permission from any world writable files
- Add the sticky bit to all world writable directories

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

{
   l_smask='01000' a_file=() a_dir=()
   a_path=(-path "*/containers/storage/*" -o -path "*/containerd/*" -o -path
"*/kubelet/*" -o -path "*/docker/*" -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 '($1 !~ /^[[:space:]]*('"$l_exclude"')/ && $2 !~
/^('"$l_exclude2"')(\/|$)/){print $2}')
}
```

References:

1. GRID: NIX-00000362
2. NIST SP 800-53 :: SC-4

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.	●	●	●

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 list any files or directories that have no owner or no group:

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

{
   a_path=(-path "*/containers/storage/*" -o -path "*/containerd/*" -o -path
"*/kubelet/*" -o -path "*/docker/*" -o -path "/sys/*" -o -path "/snap/*" -o -
path "/boot/efi/*")
   while IFS= read -r l_mount; do
      find "$l_mount" -mount -xdev \( "${a_path[@]}" \) -prune -o \
         \( -type f -o -type d \) \( -nouser -o -nogroup \) -printf '%u %g
%p\n' 2>/dev/null
   done < <(findmnt -Dkerno fstype,target | awk '($1 !~
/^[[:space:]]*(nfs|proc|cifs|smb|vfat|iso9660|efivarfs|selinuxfs|ncpfs)/ &&
$2 !~ /^(\/run|\/tmp|\/var\/tmp)(\/|$)/){print $2}')
}
```

Nothing should be returned. Each line is `<owner> <group> <path>`; a numeric owner or group indicates no matching user/group exists. Any path returned is unowned and/or ungrouped and should be investigated.

Note: on systems with many files/directories this may run for some time.

Remediation:

Remove or set ownership and group ownership of these files and/or directories to an active user on the system as appropriate.

References:

1. GRID: NIX-00000360
2. NIST SP 800-53 :: CM-6 b

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.	●	●	●

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 list all SUID and SGID files on local filesystems:

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

{
   while IFS= read -r l_mount; do
      find "$l_mount" -xdev -type f \( -perm -4000 -o -perm -2000 \) -printf
'%M %p\n' 2>/dev/null
   done < <(findmnt -Dkerno fstype,target,options | awk '($1 !~
/^[[:space:]]*(nfs|proc|smb|vfat|iso9660|efivarfs|selinuxfs)/ && $2 !~
/^\/run\/user\// && $3 !~ /noexec/ && $3 !~ /nosuid/){print $2}')
}
```

Each line is `<mode> <path>` - an `s` in the owner-execute position is SUID (`-rwsr-xr-x`), an `s` in the group-execute position is SGID (`-rwxr-sr-x`). Review each file returned and confirm it is an expected SUID/SGID program and that its integrity is intact - for example, verify the binary against its owning package with `rpm -Vf <file>`. Investigate any unexpected entry.

Note: on systems with many files this may run for some time.

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 :: AC-6, CM-6, CM-7

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.	●	●	●

7.2 Local User and Group Settings

This section provides recommendations for the integrity of the local user and group account databases: that passwords are shadowed and no account has an empty password field, that every group referenced in the user database is defined, that user and group IDs and names are unique, and that local interactive users' home directories and dot files are correctly owned and access-restricted. All accounts must have a password or be locked to prevent use by an unauthorized user. Although `useradd` and `groupadd` prevent duplicate IDs and names, manual edits to `/etc/passwd` or `/etc/group` can introduce them.

Note: The recommendations in this section check local users and groups. Any users or groups from other sources, such as LDAP, are not audited; in a domain environment, similar checks should be performed against domain users and groups.

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 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: 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
```

References:

1. GRID: NIX-00000364
2. NIST SP 800-53 :: IA-5
3. PWCONV(8)

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.		●	●

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 :: 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.		●	●

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:

Based on the Audit Procedure output, for each local interactive user whose primary group ID (the fourth field of the user's `/etc/passwd` entry) has no matching group in `/etc/group`, follow local site policy and perform one of the following:

- If the group is legitimate and should exist, create it with the reported GID:

```
# groupadd -g <GID> <groupname>
```

- If the user should instead belong to an existing group, change the user's primary group to a valid one:

```
# usermod -g <groupname> <username>
```

References:

1. GRID: NIX-00000365
2. NIST SP 800-53 :: AC-2

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			

7.2.4 Ensure no duplicate UIDs exist (Automated)

Profile Applicability:

- Level 1 - Server
- Level 1 - Workstation

Description:

Although the `useradd` program will not let you create a duplicate User ID (UID), it is possible for an administrator to manually edit the `/etc/passwd` file and change the UID field.

Rationale:

Users must be assigned unique UIDs for accountability and to ensure appropriate access protections.

Audit:

Run the following script to list any duplicate UIDs:

```
#!/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)
}
```

Nothing should be returned.

Remediation:

Based on the Audit Procedure output, for each UID shared by more than one account, keep the UID on the correct account and assign a new, unused UID to each of the other accounts:

```
# usermod -u <UID> <username>
```

`usermod` reassigns ownership of files in the account's home directory automatically. Review any other files owned by the previously shared UID and reassign ownership to the correct account as appropriate.

References:

1. GRID: NIX-00000370
2. NIST SP 800-53 :: IA-2

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			

7.2.5 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 to list any duplicate GIDs:

```
#!/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)
}
```

Nothing should be returned.

Remediation:

Based on the Audit Procedure output, for each GID shared by more than one group, keep the GID on the correct group and assign a new, unused GID to each of the other groups:

```
# groupmod -g <GID> <groupname>
```

Review any files group-owned by the previously shared GID and reassign group ownership to the correct group as appropriate.

References:

1. GRID: NIX-00000368
2. NIST SP 800-53 :: AC-2

Additional Information:

You can also use the `grpck` command to check for other inconsistencies in the `/etc/group` file.

CIS Controls:

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

7.2.6 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 username.

Rationale:

If a user is assigned a duplicate user name, it will create and have access to files with the first UID for that username in `/etc/passwd` . For example, if "test4" has a UID of 1000 and a subsequent "test4" entry has a UID of 2000, logging in as "test4" will use UID 1000. Effectively, the UID is shared, which is a security problem.

Audit:

Run the following script to list any duplicate user names:

```
#!/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)
}
```

Nothing should be returned.

Remediation:

Based on the Audit Procedure output, for each duplicated user name, determine which entry is correct and make each user name unique. Use `vipw` to safely edit the user database and either remove the redundant entry or rename it to a unique user name:

```
# vipw
```

File ownerships will reflect the change automatically as long as each account has a unique UID.

References:

1. GRID: NIX-00000371
2. NIST SP 800-53 :: IA-2

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			

7.2.7 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 to list any duplicate group names:

```
#!/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)
}
```

Nothing should be returned.

Remediation:

Based on the Audit Procedure output, for each duplicated group name, determine which entry is correct and make each group name unique. Use `vigr` to safely edit the group database and either remove the redundant entry or rename it to a unique group name:

```
# vigr
```

File group ownerships will reflect the change automatically as long as each group has a unique GID.

References:

1. GRID: NIX-00000369
2. NIST SP 800-53 :: AC-2

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			

7.2.8 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 list local interactive user home directories that do not exist, are not owned by the user, or are mode more permissive than 750:

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

{
   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 )) )"
   while IFS=" " read -r l_user l_home; do
      if [ ! -d "$l_home" ]; then
         printf '%s\n' " - User: \"$l_user\" home directory \"$l_home\" does
not exist"
      else
         while IFS=: read -r l_own l_mode; do
            [ "$l_user" != "$l_own" ] && printf '%s\n' " - User: \"$l_user\"
home \"$l_home\" is owned by \"$l_own\" (should be owned by \"$l_user\")"
            [ $(( $l_mode & $l_mask )) -gt 0 ] && printf '%s\n' " - User:
\"$l_user\" home \"$l_home\" is mode \"$l_mode\" (should be \"$l_max\" or
more restrictive)"
         done <<< "$(stat -Lc '%U:%#a' "$l_home")"
      fi
   done <<< "$(awk -v pat="$l_valid_shells" -F: '$(NF) ~ pat { print $1 " "
$(NF-1) }' /etc/passwd)"
}
```

Nothing should be returned. On a system with many local interactive users this may take some time to run.

Remediation:

1. If a local interactive user's home directory is undefined and/or does not 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 the home directory is undefined, set it to the absolute path of that directory:

```
# usermod --home <directory> <username>
```

2. For each existing home directory reported by the Audit Procedure, make changes in accordance with local site policy and apply the corresponding action:

- The directory is not owned by its user - set the owner:

```
# chown <user> <home_directory>
```

- The directory is more permissive than mode 0750 - remove the excess permissions:

```
# chmod g-w,o-rwx <home_directory>
```

References:

1. GRID: NIX-00000367
2. NIST SP 800-53 :: 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.	●	●	●

7.2.9 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.
- `.rhost` file provides the "remote authentication" database for the rcp, rlogin, and rsh commands and the rcmd() function. These files bypass the standard password-based user authentication mechanism. They specify remote hosts and users that are considered trusted (i.e. are allowed to access the local system without supplying a password)
- `.netrc` file contains data for logging into a remote host or passing authentication to an API.
- `.bash_history` file keeps track of the user's 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 list local interactive user dot files that are `.forward`/`.rhost`, or are not mode 0644 or more restrictive (0600 for `.netrc`/`.bash_history`), or are not owned by the user and their primary group:

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

{
   l_valid_shells="^($( awk -F\/ '$NF != "nologin" {print}' /etc/shells | sed
-rn '/^\//{s,/,\\/,g;p}' | paste -s -d '|' - ))$"
   f_file_chk() {
      l_max="$( printf '%o' $(( 0777 & ~$l_mask )) )"
      [ $(( $l_mode & $l_mask )) -gt 0 ] && printf '%s\n' "  - File:
\"$l_hdfile\" is mode \"$l_mode\" (should be \"$l_max\" or more restrictive)"
      [[ ! "$l_owner" =~ ($l_user) ]] && printf '%s\n' "  - File:
\"$l_hdfile\" is owned by \"$l_owner\" (should be owned by \"${l_user//|/ or
}\")"
      [[ ! "$l_gowner" =~ ($l_group) ]] && printf '%s\n' "  - File:
\"$l_hdfile\" is group owned by \"$l_gowner\" (should be group owned by
\"${l_group//|/ or }\")"
   }
   while IFS=: read -r l_user l_home; do
      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 "$l_hdfile")" in
                  .forward | .rhost ) printf '%s\n' "  - File: \"$l_hdfile\"
exists (review and manually delete)" ;;
                  .netrc | .bash_history ) l_mask='0177'; f_file_chk ;;
                  * ) l_mask='0133'; f_file_chk ;;
               esac
            done <<< "$(stat -Lc '%#a %U %G' "$l_hdfile")"
         done < <(find "$l_home" -xdev -type f -name '.*' -print0)
      fi
   done <<< "$(awk -v pat="$l_valid_shells" -F: '$(NF) ~ pat { print $1 " "
$(NF-1) }' /etc/passwd)"
}
```

Nothing should be returned. If a `.netrc` file is required and follows local site policy, it should be mode `0600` or more restrictive. On a system with many local interactive users this may take some time to run.

Remediation:

Making global modifications to users' files without alerting the user community can result in unexpected outages and unhappy users. Establish a monitoring policy to report user dot-file permissions and ownership, and take action in accordance with local site policy.

For each file reported by the Audit Procedure, in the affected user's home directory, apply the corresponding action:

- A `.forward` or `.rhosts` file exists - review its contents and remove it:

```
# rm <file>
```

- The file is more permissive than allowed - remove the excess permissions (0600 or more restrictive for `.netrc` and `.bash_history`, 0644 or more restrictive for other dot files):

```
# chmod u-x,go-rwx <file>    # .netrc and .bash_history
# chmod u-x,go-wx  <file>    # other dot files
```

- The file is not owned by its user - set the owner:

```
# chown <user> <file>
```

- The file is not group-owned by the user's primary group - set the group:

```
# chgrp <group> <file>
```

References:

1. GRID: NIX-00000366
2. NIST SP 800-53 :: 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.	●	●	●

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**		

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)	☐	☐

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 GPG keys are configured (Manual)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.2.1.2	Ensure gpgcheck is configured (Automated)	☐	☐
1.2.1.3	Ensure repo_gpgcheck is configured (Automated)	☐	☐
1.2.1.4	Ensure package manager repositories are configured (Manual)	☐	☐
1.2.1.5	Ensure weak dependencies are configured (Automated)	☐	☐
1.2.2	**Configure Package Updates**		
1.2.2.1	Ensure updates, patches, and additional security software are installed (Manual)	☐	☐
1.3	**Mandatory Access Control**		
1.3.1	**Configure SELinux**		
1.3.1.1	Ensure SELinux is installed (Automated)	☐	☐
1.3.1.2	Ensure SELinux is not disabled in bootloader configuration (Automated)	☐	☐
1.3.1.3	Ensure SELinux policy is configured (Automated)	☐	☐
1.3.1.4	Ensure the SELinux mode is not disabled (Automated)	☐	☐
1.3.1.5	Ensure the SELinux mode is enforcing (Automated)	☐	☐
1.3.1.6	Ensure no unconfined services exist (Manual)	☐	☐
1.3.1.7	Ensure the MCS Translation Service (mcstrans) is not installed (Automated)	☐	☐
1.3.1.8	Ensure SETroubleshoot is not installed (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)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.4.3	Ensure authentication is required when booting into rescue mode (Automated)	☐	☐
1.4.4	Ensure authentication is required when booting into emergency mode (Automated)	☐	☐
1.4.5	Ensure systemd ability to spawn an interactive boot process is disabled (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 fs.suid_dumpable is configured (Automated)	☐	☐
1.5.4	Ensure kernel.dmesg_restrict is configured (Automated)	☐	☐
1.5.5	Ensure kernel.kptr_restrict is configured (Automated)	☐	☐
1.5.6	Ensure kernel.randomize_va_space is configured (Automated)	☐	☐
1.5.7	Ensure kernel.yama.ptrace_scope is configured (Automated)	☐	☐
1.5.8	Ensure core file size is configured (Automated)	☐	☐
1.5.9	Ensure systemd-coredump ProcessSizeMax is configured (Automated)	☐	☐
1.5.10	Ensure systemd-coredump Storage is configured (Automated)	☐	☐
1.6	**Configure system wide crypto policy**		
1.6.1	Ensure system wide crypto policy is not legacy (Automated)	☐	☐
1.6.2	Ensure system wide crypto policy disables sha1 hash and signature support (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.6.3	Ensure system wide crypto policy macs are configured (Automated)	☐	☐
1.6.4	Ensure system wide crypto policy disables cbc for ssh (Automated)	☐	☐
1.6.5	Ensure system wide crypto policy disables chacha20-poly1305 for ssh (Manual)	☐	☐
1.6.6	Ensure system wide crypto policy disables EtM for ssh (Manual)	☐	☐
1.6.7	Ensure system wide crypto policy enables post-quantum algorithms (Automated)	☐	☐
1.7	**Configure Command Line Warning Banners**		
1.7.1	Ensure /etc/motd is configured (Automated)	☐	☐
1.7.2	Ensure /etc/issue is configured (Automated)	☐	☐
1.7.3	Ensure /etc/issue.net is configured (Automated)	☐	☐
1.7.4	Ensure access to /etc/motd is configured (Automated)	☐	☐
1.7.5	Ensure access to /etc/issue is configured (Automated)	☐	☐
1.7.6	Ensure access to /etc/issue.net is configured (Automated)	☐	☐
1.8	**Configure GNOME Display Manager**		
1.8.1	Ensure GDM login banner is configured (Automated)	☐	☐
1.8.2	Ensure GDM disable-user-list is configured (Automated)	☐	☐
1.8.3	Ensure GDM screen lock is configured (Automated)	☐	☐
1.8.4	Ensure GDM automount is configured (Automated)	☐	☐
1.8.5	Ensure GDM autorun-never is configured (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
1.8.6	Ensure XDMCP is not enabled (Automated)	☐	☐
1.8.7	Ensure Wayland 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 cockpit web services are not in use (Automated)	☐	☐
2.1.4	Ensure dhcp server services are not in use (Automated)	☐	☐
2.1.5	Ensure dns server services are not in use (Automated)	☐	☐
2.1.6	Ensure dnsmasq services are not in use (Automated)	☐	☐
2.1.7	Ensure samba file server services are not in use (Automated)	☐	☐
2.1.8	Ensure ftp server services are not in use (Automated)	☐	☐
2.1.9	Ensure message access server services are not in use (Automated)	☐	☐
2.1.10	Ensure network file system 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 snmp services are not in use (Automated)	☐	☐
2.1.15	Ensure telnet server services are not in use (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
2.1.16	Ensure tftp server services are not in use (Automated)	☐	☐
2.1.17	Ensure web proxy server services are not in use (Automated)	☐	☐
2.1.18	Ensure web server services are not in use (Automated)	☐	☐
2.1.19	Ensure mail transfer agents are configured for local-only mode (Automated)	☐	☐
2.1.20	Ensure a graphical user interface is not in use (Automated)	☐	☐
2.1.21	Ensure X window server services are not in use (Automated)	☐	☐
2.1.22	Ensure Xwayland is not in use (Automated)	☐	☐
2.1.23	Ensure only approved services are listening on a network interface (Manual)	☐	☐
2.2	**Configure Client Services**		
2.2.1	Ensure ftp client is not installed (Automated)	☐	☐
2.2.2	Ensure ldap client is not installed (Automated)	☐	☐
2.2.3	Ensure telnet client is not installed (Automated)	☐	☐
2.2.4	Ensure tftp client is not installed (Automated)	☐	☐
2.3	**Configure Time Synchronization**		
2.3.1	Ensure time synchronization is in use (Automated)	☐	☐
2.3.2	Ensure chrony is configured (Automated)	☐	☐
2.3.3	Ensure chrony is not run as the root user (Automated)	☐	☐
2.4	**Job Schedulers**		
2.4.1	**Configure cron**		

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
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)	☐	☐
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 dccp kernel module is not available (Automated)	☐	☐
3.2.2	Ensure tipc kernel module is not available (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
3.2.3	Ensure rds kernel module is not available (Automated)	☐	☐
3.2.4	Ensure sctp kernel module is not available (Automated)	☐	☐
3.2.5	Ensure atm kernel module is not available (Automated)	☐	☐
3.2.6	Ensure can kernel module is not available (Automated)	☐	☐
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)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
3.3.1.11	Ensure net.ipv4.conf.default.secure_redirects is configured (Automated)	☐	☐
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)	☐	☐
3.3.1.20	Ensure net.ipv4.conf.default.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)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
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)	☐	☐
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 firewalld**		
4.1.1	Ensure firewalld is installed (Automated)	☐	☐
4.1.2	Ensure firewalld backend is configured (Automated)	☐	☐
4.1.3	Ensure firewalld.service is configured (Automated)	☐	☐
4.1.4	Ensure firewalld active zone target is configured (Automated)	☐	☐
4.1.5	Ensure firewalld loopback traffic is configured (Manual)	☐	☐
4.1.6	Ensure firewalld loopback source address traffic is configured (Manual)	☐	☐
4.1.7	Ensure firewalld services and ports are configured (Manual)	☐	☐
4.1.8	Ensure firewalld default zone is set to drop (Manual)	☐	☐
4.1.9	Ensure conflicting firewall utilities are not in use with firewalld (Automated)	☐	☐
4.1.10	Ensure firewalld logs denied packets (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
4.1.11	Ensure firewalld blocks ICMP timestamp requests and replies (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 Ciphers are configured (Automated)	☐	☐
5.1.6	Ensure sshd KexAlgorithms is configured (Automated)	☐	☐
5.1.7	Ensure sshd MACs are configured (Automated)	☐	☐
5.1.8	Ensure sshd access is configured (Automated)	☐	☐
5.1.9	Ensure sshd Banner is configured (Automated)	☐	☐
5.1.10	Ensure sshd ClientAliveInterval and ClientAliveCountMax are configured (Automated)	☐	☐
5.1.11	Ensure sshd DisableForwarding is enabled (Automated)	☐	☐
5.1.12	Ensure sshd GSSAPIAuthentication is disabled (Automated)	☐	☐
5.1.13	Ensure sshd HostbasedAuthentication is disabled (Automated)	☐	☐
5.1.14	Ensure sshd IgnoreRhosts is enabled (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
5.1.15	Ensure sshd LoginGraceTime is configured (Automated)	☐	☐
5.1.16	Ensure sshd LogLevel is configured (Automated)	☐	☐
5.1.17	Ensure sshd MaxAuthTries is configured (Automated)	☐	☐
5.1.18	Ensure sshd MaxStartups is configured (Automated)	☐	☐
5.1.19	Ensure sshd MaxSessions is configured (Automated)	☐	☐
5.1.20	Ensure sshd PermitEmptyPasswords is disabled (Automated)	☐	☐
5.1.21	Ensure sshd PermitRootLogin is disabled (Automated)	☐	☐
5.1.22	Ensure sshd PermitUserEnvironment is disabled (Automated)	☐	☐
5.1.23	Ensure sshd UsePAM is enabled (Automated)	☐	☐
5.2	**Configure privilege escalation**		
5.2.1	Ensure sudo is installed (Automated)	☐	☐
5.2.2	Ensure access on sudoers configuration files is configured (Automated)	☐	☐
5.2.3	Ensure access on the /etc/sudoers.d directory is configured (Automated)	☐	☐
5.2.4	Ensure sudo commands use pty (Automated)	☐	☐
5.2.5	Ensure sudo log file exists (Automated)	☐	☐
5.2.6	Ensure users must provide password for escalation (Automated)	☐	☐
5.2.7	Ensure re-authentication for privilege escalation is not disabled globally (Automated)	☐	☐
5.2.8	Ensure sudo pwfeedback is configured (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
5.2.9	Ensure sudo visiblepw is configured (Automated)	☐	☐
5.2.10	Ensure sudo timestamp_timeout is configured (Automated)	☐	☐
5.2.11	Ensure sudo timestamp_type is configured (Automated)	☐	☐
5.2.12	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 authselect is installed (Automated)	☐	☐
5.3.1.3	Ensure latest version of libpwquality is installed (Automated)	☐	☐
5.3.2	**Configure authselect**		
5.3.2.1	Ensure PAM is managed by an active authselect profile (Automated)	☐	☐
5.3.2.2	Ensure pam_faillock module is enabled (Automated)	☐	☐
5.3.2.3	Ensure pam_pwquality module is enabled (Automated)	☐	☐
5.3.2.4	Ensure pam_pwhistory module is enabled (Automated)	☐	☐
5.3.2.5	Ensure pam_unix 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)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
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 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)	☐	☐
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)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
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)	☐	☐
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)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
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 in /etc/login.defs (Automated)	☐	☐
5.4.3.4	Ensure default user umask is configured (Automated)	☐	☐
6	**Logging and Auditing**		
6.1	**System Logging**		
6.1.1	**Configure journald**		
6.1.1.1	Ensure journald service is active (Automated)	☐	☐
6.1.1.2	Ensure systemd-journal-remote is not in use (Automated)	☐	☐
6.1.1.3	Ensure journald is configured to send logs to rsyslog (Automated)	☐	☐
6.1.1.4	Ensure journald log file rotation is configured (Manual)	☐	☐
6.1.1.5	Ensure journald log file access is configured (Manual)	☐	☐
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)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
6.1.2.3	Ensure rsyslog log file creation mode is configured (Automated)	☐	☐
6.1.2.4	Ensure access to rsyslog configuration files is configured (Automated)	☐	☐
6.1.2.5	Ensure access to the /etc/rsyslog.d directory is configured (Automated)	☐	☐
6.1.2.6	Ensure access to the rsyslog work directory is configured (Automated)	☐	☐
6.1.2.7	Ensure rsyslog is not configured as a log server (Automated)	☐	☐
6.1.2.8	Ensure rsyslog logging rules are configured (Manual)	☐	☐
6.1.2.9	Ensure rsyslog is configured to send logs to a remote log host (Manual)	☐	☐
6.1.2.10	Ensure a rsyslog TLS network stream driver is installed (Automated)	☐	☐
6.1.2.11	Ensure rsyslog network forwarding uses TLS (Automated)	☐	☐
6.1.2.12	Ensure rsyslog TLS CA certificate is configured (Manual)	☐	☐
6.1.2.13	Ensure rsyslog TLS client certificate is configured (Manual)	☐	☐
6.1.2.14	Ensure logrotate is 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**		

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
6.2.1.1	Ensure auditd packages are installed (Automated)	☐	☐
6.2.1.2	Ensure auditing for processes that start prior to auditd is enabled (Automated)	☐	☐
6.2.1.3	Ensure audit_backlog_limit is configured (Automated)	☐	☐
6.2.1.4	Ensure auditd service is enabled and active (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)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
6.2.3.7	Ensure events that modify /etc/hosts and /etc/hostname are collected (Automated)	☐	☐
6.2.3.8	Ensure events that modify /etc/sysconfig/network and /etc/NetworkManager/system-connections/ 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 are collected (Automated)	☐	☐
6.2.3.19	Ensure discretionary access control permission modification events chown,fchown,lchown,fchownat are collected (Automated)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
6.2.3.20	Ensure discretionary access control permission modification events setxattr,lsetxattr,fsetxattr,removexattr,lremovexattr,fremovexattr collected (Automated)	☐	☐
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)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
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)	☐	☐
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 (Automated)	☐	☐
6.2.4	**Configure auditd File Access**		
6.2.4.1	Ensure the audit log file directory mode is configured (Automated)	☐	☐
6.2.4.2	Ensure audit log files mode is configured (Automated)	☐	☐
6.2.4.3	Ensure audit log files owner is configured (Automated)	☐	☐
6.2.4.4	Ensure audit log files group owner 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)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
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)	☐	☐
6.3.4	Ensure report_url is configured (Automated)	☐	☐
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)	☐	☐

CIS Benchmark Recommendation		Set Correctly	
		Yes	No
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)	☐	☐
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 no duplicate UIDs exist (Automated)	☐	☐
7.2.5	Ensure no duplicate GIDs exist (Automated)	☐	☐
7.2.6	Ensure no duplicate user names exist (Automated)	☐	☐
7.2.7	Ensure no duplicate group names exist (Automated)	☐	☐
7.2.8	Ensure local interactive user home directories are configured (Automated)	☐	☐
7.2.9	Ensure local interactive user dot files access is configured (Automated)	☐	☐

Appendix: CIS Controls v7 IG 1 Mapped Recommendations

Recommendation		Set Correctly	
		Yes	No
1.1.2.1.1	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 GPG keys are configured	☐	☐
1.2.1.2	Ensure gpgcheck is configured	☐	☐
1.2.1.3	Ensure repo_gpgcheck is configured	☐	☐
1.2.1.4	Ensure package manager repositories are configured	☐	☐

Recommendation		Set Correctly	
		Yes	No
1.2.2.1	Ensure updates, patches, and additional security software are installed	☐	☐
1.3.1.1	Ensure SELinux is installed	☐	☐
1.3.1.2	Ensure SELinux is not disabled in bootloader configuration	☐	☐
1.3.1.3	Ensure SELinux policy is configured	☐	☐
1.3.1.4	Ensure the SELinux mode is not disabled	☐	☐
1.3.1.5	Ensure the SELinux mode is enforcing	☐	☐
1.3.1.6	Ensure no unconfined services exist	☐	☐
1.3.1.8	Ensure SETroubleshoot is not installed	☐	☐
1.4.1	Ensure bootloader password is set	☐	☐
1.4.2	Ensure access to bootloader config is configured	☐	☐
1.5.4	Ensure kernel.dmesg_restrict is configured	☐	☐
1.7.4	Ensure access to /etc/motd is configured	☐	☐
1.7.5	Ensure access to /etc/issue is configured	☐	☐
1.7.6	Ensure access to /etc/issue.net is configured	☐	☐
1.8.3	Ensure GDM screen lock is configured	☐	☐
1.8.4	Ensure GDM automount is configured	☐	☐
1.8.5	Ensure GDM autorun-never is configured	☐	☐
2.1.1	Ensure autofs services are not in use	☐	☐
2.1.15	Ensure telnet server services are not in use	☐	☐
2.2.2	Ensure ldap client is not installed	☐	☐
2.2.3	Ensure telnet 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	☐	☐

Recommendation		Set Correctly	
		Yes	No
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 firewalld is installed	☐	☐
4.1.2	Ensure firewalld backend is configured	☐	☐
4.1.3	Ensure firewalld.service is configured	☐	☐
4.1.4	Ensure firewalld active zone target is configured	☐	☐
4.1.5	Ensure firewalld loopback traffic is configured	☐	☐
4.1.6	Ensure firewalld loopback source address traffic is configured	☐	☐
4.1.7	Ensure firewalld services and ports are configured	☐	☐
5.1.1	Ensure access to /etc/ssh/sshd_config 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	☐	☐
5.1.8	Ensure sshd access is configured	☐	☐
5.1.16	Ensure sshd LogLevel is configured	☐	☐
5.1.21	Ensure sshd PermitRootLogin is disabled	☐	☐
5.2.1	Ensure sudo is installed	☐	☐
5.2.2	Ensure access on sudoers configuration files is configured	☐	☐
5.2.3	Ensure access on the /etc/sudoers.d directory is configured	☐	☐
5.2.4	Ensure sudo commands use pty	☐	☐
5.2.6	Ensure users must provide password for escalation	☐	☐
5.2.7	Ensure re-authentication for privilege escalation is not disabled globally	☐	☐
5.2.8	Ensure sudo pwfeedback is configured	☐	☐
5.2.9	Ensure sudo visiblepw is configured	☐	☐
5.2.10	Ensure sudo timestamp_timeout is configured	☐	☐
5.2.11	Ensure sudo timestamp_type is configured	☐	☐
5.2.12	Ensure access to the su command is restricted	☐	☐
5.4.1.5	Ensure inactive password lock is configured	☐	☐
5.4.2.1	Ensure root is the only UID 0 account	☐	☐

Recommendation		Set Correctly	
		Yes	No
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 in /etc/login.defs	☐	☐
5.4.3.4	Ensure default user umask is configured	☐	☐
6.1.1.1	Ensure journald service is active	☐	☐
6.1.1.3	Ensure journald is configured to send logs to rsyslog	☐	☐
6.1.1.5	Ensure journald log file access 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 access to rsyslog configuration files is configured	☐	☐
6.1.2.5	Ensure access to the /etc/rsyslog.d directory is configured	☐	☐
6.1.2.6	Ensure access to the rsyslog work directory is configured	☐	☐
6.1.2.8	Ensure rsyslog logging rules are configured	☐	☐
6.1.2.9	Ensure rsyslog is configured to send logs to a remote log host	☐	☐
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 auditing for processes that start prior to auditd is enabled	☐	☐
6.2.1.3	Ensure audit_backlog_limit is configured	☐	☐
6.2.1.4	Ensure auditd service is enabled and active	☐	☐
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	☐	☐
6.2.3.25	Ensure rename file deletion events by users are collected	☐	☐

Recommendation		Set Correctly	
		Yes	No
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 the audit log file directory mode is configured	☐	☐
6.2.4.2	Ensure audit log files mode is configured	☐	☐
6.2.4.3	Ensure audit log files owner is configured	☐	☐
6.2.4.4	Ensure audit log files group owner 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	☐	☐

Recommendation		Set Correctly	
		Yes	No
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	☐	☐
7.2.8	Ensure local interactive user home directories are configured	☐	☐
7.2.9	Ensure local interactive user dot files access is configured	☐	☐

Appendix: CIS Controls v7 IG 2 Mapped Recommendations

Recommendation		Set Correctly	
		Yes	No
1.1.1.1	Ensure 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	☐	☐

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 GPG keys are configured	☐	☐
1.2.1.2	Ensure gpgcheck is configured	☐	☐
1.2.1.3	Ensure repo_gpgcheck is configured	☐	☐
1.2.1.4	Ensure package manager repositories are configured	☐	☐
1.2.1.5	Ensure weak dependencies are configured	☐	☐
1.2.2.1	Ensure updates, patches, and additional security software are installed	☐	☐
1.3.1.1	Ensure SELinux is installed	☐	☐
1.3.1.2	Ensure SELinux is not disabled in bootloader configuration	☐	☐
1.3.1.3	Ensure SELinux policy is configured	☐	☐
1.3.1.4	Ensure the SELinux mode is not disabled	☐	☐
1.3.1.5	Ensure the SELinux mode is enforcing	☐	☐
1.3.1.6	Ensure no unconfined services exist	☐	☐
1.3.1.7	Ensure the MCS Translation Service (mcstrans) is not installed	☐	☐
1.3.1.8	Ensure SETroubleshoot is not installed	☐	☐
1.4.1	Ensure bootloader password is set	☐	☐
1.4.2	Ensure access to bootloader config is configured	☐	☐
1.5.4	Ensure kernel.dmesg_restrict is configured	☐	☐
1.5.5	Ensure kernel.kptr_restrict is configured	☐	☐
1.5.6	Ensure kernel.randomize_va_space is configured	☐	☐

Recommendation		Set Correctly	
		Yes	No
1.5.7	Ensure kernel.yama.ptrace_scope is configured	☐	☐
1.6.1	Ensure system wide crypto policy is not legacy	☐	☐
1.6.2	Ensure system wide crypto policy disables sha1 hash and signature support	☐	☐
1.6.3	Ensure system wide crypto policy macs are configured	☐	☐
1.6.4	Ensure system wide crypto policy disables cbc for ssh	☐	☐
1.6.5	Ensure system wide crypto policy disables chacha20-poly1305 for ssh	☐	☐
1.6.6	Ensure system wide crypto policy disables EtM for ssh	☐	☐
1.6.7	Ensure system wide crypto policy enables post-quantum algorithms	☐	☐
1.7.4	Ensure access to /etc/motd is configured	☐	☐
1.7.5	Ensure access to /etc/issue is configured	☐	☐
1.7.6	Ensure access to /etc/issue.net is configured	☐	☐
1.8.2	Ensure GDM disable-user-list is configured	☐	☐
1.8.3	Ensure GDM screen lock is configured	☐	☐
1.8.4	Ensure GDM automount is configured	☐	☐
1.8.5	Ensure GDM autorun-never is configured	☐	☐
1.8.6	Ensure XDMCP is not enabled	☐	☐
1.8.7	Ensure Wayland 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 cockpit web services are not in use	☐	☐
2.1.4	Ensure dhcp server services are not in use	☐	☐
2.1.5	Ensure dns server services are not in use	☐	☐
2.1.6	Ensure dnsmasq services are not in use	☐	☐
2.1.7	Ensure samba file server services are not in use	☐	☐
2.1.8	Ensure ftp server services are not in use	☐	☐
2.1.9	Ensure message access server services are not in use	☐	☐
2.1.10	Ensure network file system 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	☐	☐

Recommendation		Set Correctly	
		Yes	No
2.1.13	Ensure rsync services are not in use	☐	☐
2.1.14	Ensure snmp services are not in use	☐	☐
2.1.15	Ensure telnet server services are not in use	☐	☐
2.1.16	Ensure tftp server services are not in use	☐	☐
2.1.17	Ensure web proxy server services are not in use	☐	☐
2.1.18	Ensure web server services are not in use	☐	☐
2.1.19	Ensure mail transfer agents are configured for local-only mode	☐	☐
2.1.20	Ensure a graphical user interface is not in use	☐	☐
2.1.21	Ensure X window server services are not in use	☐	☐
2.1.22	Ensure Xwayland is not in use	☐	☐
2.1.23	Ensure only approved services are listening on a network interface	☐	☐
2.2.1	Ensure ftp client is not installed	☐	☐
2.2.2	Ensure ldap client is not installed	☐	☐
2.2.3	Ensure telnet client is not installed	☐	☐
2.2.4	Ensure tftp client is not installed	☐	☐
2.3.1	Ensure time synchronization is in use	☐	☐
2.3.2	Ensure chrony is configured	☐	☐
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 dccp kernel module is not available	☐	☐
3.2.2	Ensure tipc kernel module is not available	☐	☐

Recommendation		Set Correctly	
		Yes	No
3.2.3	Ensure rds kernel module is not available	☐	☐
3.2.4	Ensure sctp kernel module is not available	☐	☐
3.2.5	Ensure atm kernel module is not available	☐	☐
3.2.6	Ensure can kernel module is 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	☐	☐
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	☐	☐

Recommendation		Set Correctly	
		Yes	No
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 firewalld is installed	☐	☐
4.1.2	Ensure firewalld backend is configured	☐	☐
4.1.3	Ensure firewalld.service is configured	☐	☐
4.1.4	Ensure firewalld active zone target is configured	☐	☐
4.1.5	Ensure firewalld loopback traffic is configured	☐	☐
4.1.6	Ensure firewalld loopback source address traffic is configured	☐	☐
4.1.7	Ensure firewalld services and ports are configured	☐	☐
5.1.1	Ensure access to /etc/ssh/sshd_config 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 Ciphers are configured	☐	☐
5.1.6	Ensure sshd KexAlgorithms is configured	☐	☐
5.1.7	Ensure sshd MACs are configured	☐	☐
5.1.8	Ensure sshd access is configured	☐	☐
5.1.11	Ensure sshd DisableForwarding is enabled	☐	☐
5.1.12	Ensure sshd GSSAPIAuthentication is disabled	☐	☐
5.1.13	Ensure sshd HostbasedAuthentication is disabled	☐	☐
5.1.14	Ensure sshd IgnoreRhosts is enabled	☐	☐
5.1.16	Ensure sshd LogLevel is configured	☐	☐
5.1.20	Ensure sshd PermitEmptyPasswords is disabled	☐	☐
5.1.21	Ensure sshd PermitRootLogin is disabled	☐	☐
5.1.23	Ensure sshd UsePAM is enabled	☐	☐
5.2.1	Ensure sudo is installed	☐	☐

Recommendation		Set Correctly	
		Yes	No
5.2.2	Ensure access on sudoers configuration files is configured	☐	☐
5.2.3	Ensure access on the /etc/sudoers.d directory is configured	☐	☐
5.2.4	Ensure sudo commands use pty	☐	☐
5.2.5	Ensure sudo log file exists	☐	☐
5.2.6	Ensure users must provide password for escalation	☐	☐
5.2.7	Ensure re-authentication for privilege escalation is not disabled globally	☐	☐
5.2.8	Ensure sudo pwfeedback is configured	☐	☐
5.2.9	Ensure sudo visiblepw is configured	☐	☐
5.2.10	Ensure sudo timestamp_timeout is configured	☐	☐
5.2.11	Ensure sudo timestamp_type is configured	☐	☐
5.2.12	Ensure access to the su command is restricted	☐	☐
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 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	☐	☐

Recommendation		Set Correctly	
		Yes	No
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	☐	☐
5.4.1.4	Ensure strong password hashing algorithm is configured	☐	☐
5.4.1.5	Ensure inactive password lock is configured	☐	☐
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 in /etc/login.defs	☐	☐
5.4.3.4	Ensure default user umask is configured	☐	☐
6.1.1.1	Ensure journald service is active	☐	☐
6.1.1.2	Ensure systemd-journal-remote is not in use	☐	☐
6.1.1.3	Ensure journald is configured to send logs to rsyslog	☐	☐
6.1.1.4	Ensure journald log file rotation is configured	☐	☐
6.1.1.5	Ensure journald log file access 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 access to rsyslog configuration files is configured	☐	☐
6.1.2.5	Ensure access to the /etc/rsyslog.d directory is configured	☐	☐
6.1.2.6	Ensure access to the rsyslog work directory is configured	☐	☐
6.1.2.7	Ensure rsyslog is not configured as a log server	☐	☐
6.1.2.8	Ensure rsyslog logging rules are configured	☐	☐

Recommendation		Set Correctly	
		Yes	No
6.1.2.9	Ensure rsyslog is configured to send logs to a remote log host	☐	☐
6.1.2.13	Ensure rsyslog TLS client certificate is configured	☐	☐
6.1.2.14	Ensure logrotate is 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 auditing for processes that start prior to auditd is enabled	☐	☐
6.2.1.3	Ensure audit_backlog_limit is configured	☐	☐
6.2.1.4	Ensure auditd service is enabled and active	☐	☐
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	☐	☐
6.2.3.1	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 /etc/sysconfig/network and /etc/NetworkManager/system-connections/ 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	☐	☐

Recommendation		Set Correctly	
		Yes	No
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 are collected	☐	☐
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	☐	☐

Recommendation		Set Correctly	
		Yes	No
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 the audit log file directory mode is configured	☐	☐
6.2.4.2	Ensure audit log files mode is configured	☐	☐
6.2.4.3	Ensure audit log files owner is configured	☐	☐
6.2.4.4	Ensure audit log files group owner 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	☐	☐

Recommendation		Set Correctly	
		Yes	No
7.2.1	Ensure accounts in /etc/passwd use shadowed passwords	☐	☐
7.2.2	Ensure /etc/shadow password fields are not empty	☐	☐
7.2.8	Ensure local interactive user home directories are configured	☐	☐
7.2.9	Ensure local interactive user dot files access is configured	☐	☐

Appendix: CIS Controls v7 IG 3 Mapped Recommendations

Recommendation		Set Correctly	
		Yes	No
1.1.1.1	Ensure 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	☐	☐

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 GPG keys are configured	☐	☐
1.2.1.2	Ensure gpgcheck is configured	☐	☐
1.2.1.3	Ensure repo_gpgcheck is configured	☐	☐
1.2.1.4	Ensure package manager repositories are configured	☐	☐
1.2.1.5	Ensure weak dependencies are configured	☐	☐
1.2.2.1	Ensure updates, patches, and additional security software are installed	☐	☐
1.3.1.1	Ensure SELinux is installed	☐	☐
1.3.1.2	Ensure SELinux is not disabled in bootloader configuration	☐	☐
1.3.1.3	Ensure SELinux policy is configured	☐	☐
1.3.1.4	Ensure the SELinux mode is not disabled	☐	☐
1.3.1.5	Ensure the SELinux mode is enforcing	☐	☐
1.3.1.6	Ensure no unconfined services exist	☐	☐
1.3.1.7	Ensure the MCS Translation Service (mcstrans) is not installed	☐	☐
1.3.1.8	Ensure SETroubleshoot is not installed	☐	☐
1.4.1	Ensure bootloader password is set	☐	☐
1.4.2	Ensure access to bootloader config is configured	☐	☐
1.5.4	Ensure kernel.dmesg_restrict is configured	☐	☐
1.5.5	Ensure kernel.kptr_restrict is configured	☐	☐
1.5.6	Ensure kernel.randomize_va_space is configured	☐	☐

Recommendation		Set Correctly	
		Yes	No
1.5.7	Ensure kernel.yama.ptrace_scope is configured	☐	☐
1.6.1	Ensure system wide crypto policy is not legacy	☐	☐
1.6.2	Ensure system wide crypto policy disables sha1 hash and signature support	☐	☐
1.6.3	Ensure system wide crypto policy macs are configured	☐	☐
1.6.4	Ensure system wide crypto policy disables cbc for ssh	☐	☐
1.6.5	Ensure system wide crypto policy disables chacha20-poly1305 for ssh	☐	☐
1.6.6	Ensure system wide crypto policy disables EtM for ssh	☐	☐
1.6.7	Ensure system wide crypto policy enables post-quantum algorithms	☐	☐
1.7.4	Ensure access to /etc/motd is configured	☐	☐
1.7.5	Ensure access to /etc/issue is configured	☐	☐
1.7.6	Ensure access to /etc/issue.net is configured	☐	☐
1.8.2	Ensure GDM disable-user-list is configured	☐	☐
1.8.3	Ensure GDM screen lock is configured	☐	☐
1.8.4	Ensure GDM automount is configured	☐	☐
1.8.5	Ensure GDM autorun-never is configured	☐	☐
1.8.6	Ensure XDMCP is not enabled	☐	☐
1.8.7	Ensure Wayland 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 cockpit web services are not in use	☐	☐
2.1.4	Ensure dhcp server services are not in use	☐	☐
2.1.5	Ensure dns server services are not in use	☐	☐
2.1.6	Ensure dnsmasq services are not in use	☐	☐
2.1.7	Ensure samba file server services are not in use	☐	☐
2.1.8	Ensure ftp server services are not in use	☐	☐
2.1.9	Ensure message access server services are not in use	☐	☐
2.1.10	Ensure network file system 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	☐	☐

Recommendation		Set Correctly	
		Yes	No
2.1.13	Ensure rsync services are not in use	☐	☐
2.1.14	Ensure snmp services are not in use	☐	☐
2.1.15	Ensure telnet server services are not in use	☐	☐
2.1.16	Ensure tftp server services are not in use	☐	☐
2.1.17	Ensure web proxy server services are not in use	☐	☐
2.1.18	Ensure web server services are not in use	☐	☐
2.1.19	Ensure mail transfer agents are configured for local-only mode	☐	☐
2.1.20	Ensure a graphical user interface is not in use	☐	☐
2.1.21	Ensure X window server services are not in use	☐	☐
2.1.22	Ensure Xwayland is not in use	☐	☐
2.1.23	Ensure only approved services are listening on a network interface	☐	☐
2.2.1	Ensure ftp client is not installed	☐	☐
2.2.2	Ensure ldap client is not installed	☐	☐
2.2.3	Ensure telnet client is not installed	☐	☐
2.2.4	Ensure tftp client is not installed	☐	☐
2.3.1	Ensure time synchronization is in use	☐	☐
2.3.2	Ensure chrony is configured	☐	☐
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 dccp kernel module is not available	☐	☐

Recommendation		Set Correctly	
		Yes	No
3.2.2	Ensure tipc kernel module is not available	☐	☐
3.2.3	Ensure rds kernel module is not available	☐	☐
3.2.4	Ensure sctp kernel module is not available	☐	☐
3.2.5	Ensure atm kernel module is not available	☐	☐
3.2.6	Ensure can kernel module is 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	☐	☐
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	☐	☐

Recommendation		Set Correctly	
		Yes	No
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 firewalld is installed	☐	☐
4.1.2	Ensure firewalld backend is configured	☐	☐
4.1.3	Ensure firewalld.service is configured	☐	☐
4.1.4	Ensure firewalld active zone target is configured	☐	☐
4.1.5	Ensure firewalld loopback traffic is configured	☐	☐
4.1.6	Ensure firewalld loopback source address traffic is configured	☐	☐
4.1.7	Ensure firewalld services and ports are configured	☐	☐
5.1.1	Ensure access to /etc/ssh/sshd_config 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 Ciphers are configured	☐	☐
5.1.6	Ensure sshd KexAlgorithms is configured	☐	☐
5.1.7	Ensure sshd MACs are configured	☐	☐
5.1.8	Ensure sshd access is configured	☐	☐
5.1.11	Ensure sshd DisableForwarding is enabled	☐	☐
5.1.12	Ensure sshd GSSAPIAuthentication is disabled	☐	☐
5.1.13	Ensure sshd HostbasedAuthentication is disabled	☐	☐
5.1.14	Ensure sshd IgnoreRhosts is enabled	☐	☐
5.1.16	Ensure sshd LogLevel is configured	☐	☐
5.1.17	Ensure sshd MaxAuthTries is configured	☐	☐
5.1.20	Ensure sshd PermitEmptyPasswords is disabled	☐	☐
5.1.21	Ensure sshd PermitRootLogin is disabled	☐	☐
5.1.23	Ensure sshd UsePAM is enabled	☐	☐

Recommendation		Set Correctly	
		Yes	No
5.2.1	Ensure sudo is installed	☐	☐
5.2.2	Ensure access on sudoers configuration files is configured	☐	☐
5.2.3	Ensure access on the /etc/sudoers.d directory is configured	☐	☐
5.2.4	Ensure sudo commands use pty	☐	☐
5.2.5	Ensure sudo log file exists	☐	☐
5.2.6	Ensure users must provide password for escalation	☐	☐
5.2.7	Ensure re-authentication for privilege escalation is not disabled globally	☐	☐
5.2.8	Ensure sudo pwfeedback is configured	☐	☐
5.2.9	Ensure sudo visiblepw is configured	☐	☐
5.2.10	Ensure sudo timestamp_timeout is configured	☐	☐
5.2.11	Ensure sudo timestamp_type is configured	☐	☐
5.2.12	Ensure access to the su command is restricted	☐	☐
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 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	☐	☐

Recommendation		Set Correctly	
		Yes	No
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	☐	☐
5.4.1.4	Ensure strong password hashing algorithm is configured	☐	☐
5.4.1.5	Ensure inactive password lock is configured	☐	☐
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 in /etc/login.defs	☐	☐
5.4.3.4	Ensure default user umask is configured	☐	☐
6.1.1.1	Ensure journald service is active	☐	☐
6.1.1.2	Ensure systemd-journal-remote is not in use	☐	☐
6.1.1.3	Ensure journald is configured to send logs to rsyslog	☐	☐
6.1.1.4	Ensure journald log file rotation is configured	☐	☐
6.1.1.5	Ensure journald log file access 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 access to rsyslog configuration files is configured	☐	☐
6.1.2.5	Ensure access to the /etc/rsyslog.d directory is configured	☐	☐
6.1.2.6	Ensure access to the rsyslog work directory is configured	☐	☐
6.1.2.7	Ensure rsyslog is not configured as a log server	☐	☐
6.1.2.8	Ensure rsyslog logging rules are configured	☐	☐

Recommendation		Set Correctly	
		Yes	No
6.1.2.9	Ensure rsyslog is configured to send logs to a remote log host	☐	☐
6.1.2.13	Ensure rsyslog TLS client certificate is configured	☐	☐
6.1.2.14	Ensure logrotate is 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 auditing for processes that start prior to auditd is enabled	☐	☐
6.2.1.3	Ensure audit_backlog_limit is configured	☐	☐
6.2.1.4	Ensure auditd service is enabled and active	☐	☐
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	☐	☐
6.2.3.1	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 /etc/sysconfig/network and /etc/NetworkManager/system-connections/ 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	☐	☐

Recommendation		Set Correctly	
		Yes	No
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 are collected	☐	☐
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	☐	☐

Recommendation		Set Correctly	
		Yes	No
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 the audit log file directory mode is configured	☐	☐
6.2.4.2	Ensure audit log files mode is configured	☐	☐
6.2.4.3	Ensure audit log files owner is configured	☐	☐
6.2.4.4	Ensure audit log files group owner 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	☐	☐
6.3.3	Ensure cryptographic mechanisms are used to protect the integrity of audit tools	☐	☐
6.3.4	Ensure report_url 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	☐	☐

Recommendation		Set Correctly	
		Yes	No
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	☐	☐
7.2.8	Ensure local interactive user home directories are configured	☐	☐
7.2.9	Ensure local interactive user dot files access is configured	☐	☐

Appendix: CIS Controls v7 Unmapped Recommendations

Recommendation		Set Correctly	
		Yes	No
1.1.2.5.1	Ensure separate partition exists for /var/tmp	☐	☐
1.4.3	Ensure authentication is required when booting into rescue mode	☐	☐
1.4.4	Ensure authentication is required when booting into emergency mode	☐	☐
1.4.5	Ensure systemd ability to spawn an interactive boot process is disabled	☐	☐
3.3.1.20	Ensure net.ipv4.conf.default.route_localnet is configured	☐	☐
4.1.8	Ensure firewalld default zone is set to drop	☐	☐
4.1.9	Ensure conflicting firewall utilities are not in use with firewalld	☐	☐
4.1.10	Ensure firewalld logs denied packets	☐	☐
4.1.11	Ensure firewalld blocks ICMP timestamp requests and replies	☐	☐
5.1.2	Ensure access to /etc/ssh/sshd_config.d/ directory is configured	☐	☐
6.1.2.10	Ensure a rsyslog TLS network stream driver is installed	☐	☐
6.1.2.11	Ensure rsyslog network forwarding uses TLS	☐	☐
6.1.2.12	Ensure rsyslog TLS CA certificate is configured	☐	☐

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	☐	☐
1.2.1.1	Ensure GPG keys are configured	☐	☐

Recommendation		Set Correctly	
		Yes	No
1.2.1.2	Ensure gpgcheck is configured	☐	☐
1.2.1.3	Ensure repo_gpgcheck is configured	☐	☐
1.2.1.4	Ensure package manager repositories are configured	☐	☐
1.2.2.1	Ensure updates, patches, and additional security software are installed	☐	☐
1.3.1.1	Ensure SELinux is installed	☐	☐
1.3.1.2	Ensure SELinux is not disabled in bootloader configuration	☐	☐
1.3.1.3	Ensure SELinux policy is configured	☐	☐
1.3.1.4	Ensure the SELinux mode is not disabled	☐	☐
1.3.1.5	Ensure the SELinux mode is enforcing	☐	☐
1.3.1.6	Ensure no unconfined services exist	☐	☐
1.4.1	Ensure bootloader password is set	☐	☐
1.4.2	Ensure access to bootloader config is configured	☐	☐
1.5.4	Ensure kernel.dmesg_restrict is configured	☐	☐
1.7.4	Ensure access to /etc/motd is configured	☐	☐
1.7.5	Ensure access to /etc/issue is configured	☐	☐
1.7.6	Ensure access to /etc/issue.net is configured	☐	☐
1.8.2	Ensure GDM disable-user-list is configured	☐	☐
1.8.3	Ensure GDM screen lock is configured	☐	☐
1.8.4	Ensure GDM automount is configured	☐	☐
1.8.5	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	☐	☐

Recommendation		Set Correctly	
		Yes	No
4.1.1	Ensure firewalld is installed	☐	☐
4.1.2	Ensure firewalld backend is configured	☐	☐
4.1.3	Ensure firewalld.service is configured	☐	☐
4.1.4	Ensure firewalld active zone target is configured	☐	☐
4.1.5	Ensure firewalld loopback traffic is configured	☐	☐
4.1.6	Ensure firewalld loopback source address traffic is configured	☐	☐
4.1.7	Ensure firewalld services and ports are configured	☐	☐
5.1.1	Ensure access to /etc/ssh/sshd_config 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	☐	☐
5.1.8	Ensure sshd access is configured	☐	☐
5.1.12	Ensure sshd GSSAPIAuthentication is disabled	☐	☐
5.1.13	Ensure sshd HostbasedAuthentication is disabled	☐	☐
5.1.14	Ensure sshd IgnoreRhosts is enabled	☐	☐
5.1.16	Ensure sshd LogLevel is configured	☐	☐
5.1.20	Ensure sshd PermitEmptyPasswords is disabled	☐	☐
5.1.21	Ensure sshd PermitRootLogin is disabled	☐	☐
5.1.23	Ensure sshd UsePAM is enabled	☐	☐
5.2.1	Ensure sudo is installed	☐	☐
5.2.2	Ensure access on sudoers configuration files is configured	☐	☐
5.2.3	Ensure access on the /etc/sudoers.d directory is configured	☐	☐
5.2.4	Ensure sudo commands use pty	☐	☐
5.2.6	Ensure users must provide password for escalation	☐	☐
5.2.7	Ensure re-authentication for privilege escalation is not disabled globally	☐	☐
5.2.8	Ensure sudo pwfeedback is configured	☐	☐
5.2.9	Ensure sudo visiblepw is configured	☐	☐
5.2.10	Ensure sudo timestamp_timeout is configured	☐	☐
5.2.11	Ensure sudo timestamp_type is configured	☐	☐

Recommendation		Set Correctly	
		Yes	No
5.2.12	Ensure access to the su command is restricted	☐	☐
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 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	☐	☐
5.3.3.4.4	Ensure pam_unix includes use_authtok	☐	☐
5.4.1.5	Ensure inactive password lock is configured	☐	☐
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	☐	☐

Recommendation		Set Correctly	
		Yes	No
5.4.3.3	Ensure default user umask is configured in /etc/login.defs	☐	☐
5.4.3.4	Ensure default user umask is configured	☐	☐
6.1.1.1	Ensure journald service is active	☐	☐
6.1.1.3	Ensure journald is configured to send logs to rsyslog	☐	☐
6.1.1.4	Ensure journald log file rotation is configured	☐	☐
6.1.1.5	Ensure journald log file access 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 access to rsyslog configuration files is configured	☐	☐
6.1.2.5	Ensure access to the /etc/rsyslog.d directory is configured	☐	☐
6.1.2.6	Ensure access to the rsyslog work directory is configured	☐	☐
6.1.2.8	Ensure rsyslog logging rules are configured	☐	☐
6.1.2.9	Ensure rsyslog is configured to send logs to a remote log host	☐	☐
6.1.2.14	Ensure logrotate is 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 auditing for processes that start prior to auditd is enabled	☐	☐
6.2.1.3	Ensure audit_backlog_limit is configured	☐	☐
6.2.1.4	Ensure auditd service is enabled and active	☐	☐
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	☐	☐

Recommendation		Set Correctly	
		Yes	No
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 the audit log file directory mode is configured	☐	☐
6.2.4.2	Ensure audit log files mode is configured	☐	☐
6.2.4.3	Ensure audit log files owner is configured	☐	☐
6.2.4.4	Ensure audit log files group owner 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	☐	☐
7.2.2	Ensure /etc/shadow password fields are not empty	☐	☐
7.2.8	Ensure local interactive user home directories are configured	☐	☐
7.2.9	Ensure local interactive user dot files access is configured	☐	☐

Appendix: CIS Controls v8 IG 2 Mapped Recommendations

Recommendation		Set Correctly	
		Yes	No
1.1.1.1	Ensure 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	☐	☐

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 GPG keys are configured	☐	☐
1.2.1.2	Ensure gpgcheck is configured	☐	☐
1.2.1.3	Ensure repo_gpgcheck is configured	☐	☐
1.2.1.4	Ensure package manager repositories are configured	☐	☐
1.2.1.5	Ensure weak dependencies are configured	☐	☐
1.2.2.1	Ensure updates, patches, and additional security software are installed	☐	☐
1.3.1.1	Ensure SELinux is installed	☐	☐
1.3.1.2	Ensure SELinux is not disabled in bootloader configuration	☐	☐
1.3.1.3	Ensure SELinux policy is configured	☐	☐
1.3.1.4	Ensure the SELinux mode is not disabled	☐	☐
1.3.1.5	Ensure the SELinux mode is enforcing	☐	☐
1.3.1.6	Ensure no unconfined services exist	☐	☐
1.3.1.7	Ensure the MCS Translation Service (mcstrans) is not installed	☐	☐
1.3.1.8	Ensure SETroubleshoot is not installed	☐	☐
1.4.1	Ensure bootloader password is set	☐	☐
1.4.2	Ensure access to bootloader config is configured	☐	☐
1.5.4	Ensure kernel.dmesg_restrict is configured	☐	☐
1.5.6	Ensure kernel.randomize_va_space is configured	☐	☐
1.5.7	Ensure kernel.yama.ptrace_scope is configured	☐	☐

Recommendation		Set Correctly	
		Yes	No
1.6.1	Ensure system wide crypto policy is not legacy	☐	☐
1.6.2	Ensure system wide crypto policy disables sha1 hash and signature support	☐	☐
1.6.3	Ensure system wide crypto policy macs are configured	☐	☐
1.6.4	Ensure system wide crypto policy disables cbc for ssh	☐	☐
1.6.5	Ensure system wide crypto policy disables chacha20-poly1305 for ssh	☐	☐
1.6.6	Ensure system wide crypto policy disables EtM for ssh	☐	☐
1.6.7	Ensure system wide crypto policy enables post-quantum algorithms	☐	☐
1.7.4	Ensure access to /etc/motd is configured	☐	☐
1.7.5	Ensure access to /etc/issue is configured	☐	☐
1.7.6	Ensure access to /etc/issue.net is configured	☐	☐
1.8.2	Ensure GDM disable-user-list is configured	☐	☐
1.8.3	Ensure GDM screen lock is configured	☐	☐
1.8.4	Ensure GDM automount is configured	☐	☐
1.8.5	Ensure GDM autorun-never is configured	☐	☐
1.8.6	Ensure XDMCP is not enabled	☐	☐
1.8.7	Ensure Wayland 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 cockpit web services are not in use	☐	☐
2.1.4	Ensure dhcp server services are not in use	☐	☐
2.1.5	Ensure dns server services are not in use	☐	☐
2.1.6	Ensure dnsmasq services are not in use	☐	☐
2.1.7	Ensure samba file server services are not in use	☐	☐
2.1.8	Ensure ftp server services are not in use	☐	☐
2.1.9	Ensure message access server services are not in use	☐	☐
2.1.10	Ensure network file system 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	☐	☐

Recommendation		Set Correctly	
		Yes	No
2.1.14	Ensure snmp services are not in use	☐	☐
2.1.15	Ensure telnet server services are not in use	☐	☐
2.1.16	Ensure tftp server services are not in use	☐	☐
2.1.17	Ensure web proxy server services are not in use	☐	☐
2.1.18	Ensure web server services are not in use	☐	☐
2.1.19	Ensure mail transfer agents are configured for local-only mode	☐	☐
2.1.20	Ensure a graphical user interface is not in use	☐	☐
2.1.21	Ensure X window server services are not in use	☐	☐
2.1.22	Ensure Xwayland is not in use	☐	☐
2.1.23	Ensure only approved services are listening on a network interface	☐	☐
2.2.1	Ensure ftp client is not installed	☐	☐
2.2.2	Ensure ldap client is not installed	☐	☐
2.2.3	Ensure telnet client is not installed	☐	☐
2.2.4	Ensure tftp client is not installed	☐	☐
2.3.1	Ensure time synchronization is in use	☐	☐
2.3.2	Ensure chrony is configured	☐	☐
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 dccp kernel module is not available	☐	☐
3.2.2	Ensure tipc kernel module is not available	☐	☐

Recommendation		Set Correctly	
		Yes	No
3.2.3	Ensure rds kernel module is not available	☐	☐
3.2.4	Ensure sctp kernel module is not available	☐	☐
3.2.5	Ensure atm kernel module is not available	☐	☐
3.2.6	Ensure can kernel module is 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	☐	☐
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	☐	☐

Recommendation		Set Correctly	
		Yes	No
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 firewalld is installed	☐	☐
4.1.2	Ensure firewalld backend is configured	☐	☐
4.1.3	Ensure firewalld.service is configured	☐	☐
4.1.4	Ensure firewalld active zone target is configured	☐	☐
4.1.5	Ensure firewalld loopback traffic is configured	☐	☐
4.1.6	Ensure firewalld loopback source address traffic is configured	☐	☐
4.1.7	Ensure firewalld services and ports are configured	☐	☐
5.1.1	Ensure access to /etc/ssh/sshd_config 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 Ciphers are configured	☐	☐
5.1.6	Ensure sshd KexAlgorithms is configured	☐	☐
5.1.7	Ensure sshd MACs are configured	☐	☐
5.1.8	Ensure sshd access is configured	☐	☐
5.1.11	Ensure sshd DisableForwarding is enabled	☐	☐
5.1.12	Ensure sshd GSSAPIAuthentication is disabled	☐	☐
5.1.13	Ensure sshd HostbasedAuthentication is disabled	☐	☐
5.1.14	Ensure sshd IgnoreRhosts is enabled	☐	☐
5.1.16	Ensure sshd LogLevel is configured	☐	☐
5.1.17	Ensure sshd MaxAuthTries is configured	☐	☐
5.1.20	Ensure sshd PermitEmptyPasswords is disabled	☐	☐
5.1.21	Ensure sshd PermitRootLogin is disabled	☐	☐
5.1.23	Ensure sshd UsePAM is enabled	☐	☐

Recommendation		Set Correctly	
		Yes	No
5.2.1	Ensure sudo is installed	☐	☐
5.2.2	Ensure access on sudoers configuration files is configured	☐	☐
5.2.3	Ensure access on the /etc/sudoers.d directory is configured	☐	☐
5.2.4	Ensure sudo commands use pty	☐	☐
5.2.5	Ensure sudo log file exists	☐	☐
5.2.6	Ensure users must provide password for escalation	☐	☐
5.2.7	Ensure re-authentication for privilege escalation is not disabled globally	☐	☐
5.2.8	Ensure sudo pwfeedback is configured	☐	☐
5.2.9	Ensure sudo visiblepw is configured	☐	☐
5.2.10	Ensure sudo timestamp_timeout is configured	☐	☐
5.2.11	Ensure sudo timestamp_type is configured	☐	☐
5.2.12	Ensure access to the su command is restricted	☐	☐
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 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	☐	☐

Recommendation		Set Correctly	
		Yes	No
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	☐	☐
5.4.1.4	Ensure strong password hashing algorithm is configured	☐	☐
5.4.1.5	Ensure inactive password lock is configured	☐	☐
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 in /etc/login.defs	☐	☐
5.4.3.4	Ensure default user umask is configured	☐	☐
6.1.1.1	Ensure journald service is active	☐	☐
6.1.1.2	Ensure systemd-journal-remote is not in use	☐	☐
6.1.1.3	Ensure journald is configured to send logs to rsyslog	☐	☐
6.1.1.4	Ensure journald log file rotation is configured	☐	☐
6.1.1.5	Ensure journald log file access 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 access to rsyslog configuration files is configured	☐	☐
6.1.2.5	Ensure access to the /etc/rsyslog.d directory is configured	☐	☐
6.1.2.6	Ensure access to the rsyslog work directory is configured	☐	☐
6.1.2.7	Ensure rsyslog is not configured as a log server	☐	☐
6.1.2.8	Ensure rsyslog logging rules are configured	☐	☐

Recommendation		Set Correctly	
		Yes	No
6.1.2.9	Ensure rsyslog is configured to send logs to a remote log host	☐	☐
6.1.2.13	Ensure rsyslog TLS client certificate is configured	☐	☐
6.1.2.14	Ensure logrotate is 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 auditing for processes that start prior to auditd is enabled	☐	☐
6.2.1.3	Ensure audit_backlog_limit is configured	☐	☐
6.2.1.4	Ensure auditd service is enabled and active	☐	☐
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	☐	☐
6.2.3.1	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 /etc/sysconfig/network and /etc/NetworkManager/system-connections/ 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	☐	☐

Recommendation		Set Correctly	
		Yes	No
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 are collected	☐	☐
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	☐	☐

Recommendation		Set Correctly	
		Yes	No
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 the audit log file directory mode is configured	☐	☐
6.2.4.2	Ensure audit log files mode is configured	☐	☐
6.2.4.3	Ensure audit log files owner is configured	☐	☐
6.2.4.4	Ensure audit log files group owner 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	☐	☐

Recommendation		Set Correctly	
		Yes	No
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	☐	☐
7.2.8	Ensure local interactive user home directories are configured	☐	☐
7.2.9	Ensure local interactive user dot files access is configured	☐	☐

Appendix: CIS Controls v8 IG 3 Mapped Recommendations

Recommendation		Set Correctly	
		Yes	No
1.1.1.1	Ensure 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	☐	☐

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 GPG keys are configured	☐	☐
1.2.1.2	Ensure gpgcheck is configured	☐	☐
1.2.1.3	Ensure repo_gpgcheck is configured	☐	☐
1.2.1.4	Ensure package manager repositories are configured	☐	☐
1.2.1.5	Ensure weak dependencies are configured	☐	☐
1.2.2.1	Ensure updates, patches, and additional security software are installed	☐	☐
1.3.1.1	Ensure SELinux is installed	☐	☐
1.3.1.2	Ensure SELinux is not disabled in bootloader configuration	☐	☐
1.3.1.3	Ensure SELinux policy is configured	☐	☐
1.3.1.4	Ensure the SELinux mode is not disabled	☐	☐
1.3.1.5	Ensure the SELinux mode is enforcing	☐	☐
1.3.1.6	Ensure no unconfined services exist	☐	☐
1.3.1.7	Ensure the MCS Translation Service (mcstrans) is not installed	☐	☐
1.3.1.8	Ensure SETroubleshoot is not installed	☐	☐
1.4.1	Ensure bootloader password is set	☐	☐
1.4.2	Ensure access to bootloader config is configured	☐	☐
1.5.4	Ensure kernel.dmesg_restrict is configured	☐	☐
1.5.6	Ensure kernel.randomize_va_space is configured	☐	☐
1.5.7	Ensure kernel.yama.ptrace_scope is configured	☐	☐

Recommendation		Set Correctly	
		Yes	No
1.6.1	Ensure system wide crypto policy is not legacy	☐	☐
1.6.2	Ensure system wide crypto policy disables sha1 hash and signature support	☐	☐
1.6.3	Ensure system wide crypto policy macs are configured	☐	☐
1.6.4	Ensure system wide crypto policy disables cbc for ssh	☐	☐
1.6.5	Ensure system wide crypto policy disables chacha20-poly1305 for ssh	☐	☐
1.6.6	Ensure system wide crypto policy disables EtM for ssh	☐	☐
1.6.7	Ensure system wide crypto policy enables post-quantum algorithms	☐	☐
1.7.4	Ensure access to /etc/motd is configured	☐	☐
1.7.5	Ensure access to /etc/issue is configured	☐	☐
1.7.6	Ensure access to /etc/issue.net is configured	☐	☐
1.8.2	Ensure GDM disable-user-list is configured	☐	☐
1.8.3	Ensure GDM screen lock is configured	☐	☐
1.8.4	Ensure GDM automount is configured	☐	☐
1.8.5	Ensure GDM autorun-never is configured	☐	☐
1.8.6	Ensure XDMCP is not enabled	☐	☐
1.8.7	Ensure Wayland 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 cockpit web services are not in use	☐	☐
2.1.4	Ensure dhcp server services are not in use	☐	☐
2.1.5	Ensure dns server services are not in use	☐	☐
2.1.6	Ensure dnsmasq services are not in use	☐	☐
2.1.7	Ensure samba file server services are not in use	☐	☐
2.1.8	Ensure ftp server services are not in use	☐	☐
2.1.9	Ensure message access server services are not in use	☐	☐
2.1.10	Ensure network file system 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	☐	☐

Recommendation		Set Correctly	
		Yes	No
2.1.14	Ensure snmp services are not in use	☐	☐
2.1.15	Ensure telnet server services are not in use	☐	☐
2.1.16	Ensure tftp server services are not in use	☐	☐
2.1.17	Ensure web proxy server services are not in use	☐	☐
2.1.18	Ensure web server services are not in use	☐	☐
2.1.19	Ensure mail transfer agents are configured for local-only mode	☐	☐
2.1.20	Ensure a graphical user interface is not in use	☐	☐
2.1.21	Ensure X window server services are not in use	☐	☐
2.1.22	Ensure Xwayland is not in use	☐	☐
2.1.23	Ensure only approved services are listening on a network interface	☐	☐
2.2.1	Ensure ftp client is not installed	☐	☐
2.2.2	Ensure ldap client is not installed	☐	☐
2.2.3	Ensure telnet client is not installed	☐	☐
2.2.4	Ensure tftp client is not installed	☐	☐
2.3.1	Ensure time synchronization is in use	☐	☐
2.3.2	Ensure chrony is configured	☐	☐
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 dccp kernel module is not available	☐	☐
3.2.2	Ensure tipc kernel module is not available	☐	☐

Recommendation		Set Correctly	
		Yes	No
3.2.3	Ensure rds kernel module is not available	☐	☐
3.2.4	Ensure sctp kernel module is not available	☐	☐
3.2.5	Ensure atm kernel module is not available	☐	☐
3.2.6	Ensure can kernel module is 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	☐	☐
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	☐	☐

Recommendation		Set Correctly	
		Yes	No
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 firewalld is installed	☐	☐
4.1.2	Ensure firewalld backend is configured	☐	☐
4.1.3	Ensure firewalld.service is configured	☐	☐
4.1.4	Ensure firewalld active zone target is configured	☐	☐
4.1.5	Ensure firewalld loopback traffic is configured	☐	☐
4.1.6	Ensure firewalld loopback source address traffic is configured	☐	☐
4.1.7	Ensure firewalld services and ports are configured	☐	☐
5.1.1	Ensure access to /etc/ssh/sshd_config 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 Ciphers are configured	☐	☐
5.1.6	Ensure sshd KexAlgorithms is configured	☐	☐
5.1.7	Ensure sshd MACs are configured	☐	☐
5.1.8	Ensure sshd access is configured	☐	☐
5.1.11	Ensure sshd DisableForwarding is enabled	☐	☐
5.1.12	Ensure sshd GSSAPIAuthentication is disabled	☐	☐
5.1.13	Ensure sshd HostbasedAuthentication is disabled	☐	☐
5.1.14	Ensure sshd IgnoreRhosts is enabled	☐	☐
5.1.16	Ensure sshd LogLevel is configured	☐	☐
5.1.17	Ensure sshd MaxAuthTries is configured	☐	☐
5.1.20	Ensure sshd PermitEmptyPasswords is disabled	☐	☐
5.1.21	Ensure sshd PermitRootLogin is disabled	☐	☐
5.1.23	Ensure sshd UsePAM is enabled	☐	☐

Recommendation		Set Correctly	
		Yes	No
5.2.1	Ensure sudo is installed	☐	☐
5.2.2	Ensure access on sudoers configuration files is configured	☐	☐
5.2.3	Ensure access on the /etc/sudoers.d directory is configured	☐	☐
5.2.4	Ensure sudo commands use pty	☐	☐
5.2.5	Ensure sudo log file exists	☐	☐
5.2.6	Ensure users must provide password for escalation	☐	☐
5.2.7	Ensure re-authentication for privilege escalation is not disabled globally	☐	☐
5.2.8	Ensure sudo pwfeedback is configured	☐	☐
5.2.9	Ensure sudo visiblepw is configured	☐	☐
5.2.10	Ensure sudo timestamp_timeout is configured	☐	☐
5.2.11	Ensure sudo timestamp_type is configured	☐	☐
5.2.12	Ensure access to the su command is restricted	☐	☐
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 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	☐	☐

Recommendation		Set Correctly	
		Yes	No
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	☐	☐
5.4.1.4	Ensure strong password hashing algorithm is configured	☐	☐
5.4.1.5	Ensure inactive password lock is configured	☐	☐
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 in /etc/login.defs	☐	☐
5.4.3.4	Ensure default user umask is configured	☐	☐
6.1.1.1	Ensure journald service is active	☐	☐
6.1.1.2	Ensure systemd-journal-remote is not in use	☐	☐
6.1.1.3	Ensure journald is configured to send logs to rsyslog	☐	☐
6.1.1.4	Ensure journald log file rotation is configured	☐	☐
6.1.1.5	Ensure journald log file access 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 access to rsyslog configuration files is configured	☐	☐
6.1.2.5	Ensure access to the /etc/rsyslog.d directory is configured	☐	☐
6.1.2.6	Ensure access to the rsyslog work directory is configured	☐	☐
6.1.2.7	Ensure rsyslog is not configured as a log server	☐	☐
6.1.2.8	Ensure rsyslog logging rules are configured	☐	☐

Recommendation		Set Correctly	
		Yes	No
6.1.2.9	Ensure rsyslog is configured to send logs to a remote log host	☐	☐
6.1.2.13	Ensure rsyslog TLS client certificate is configured	☐	☐
6.1.2.14	Ensure logrotate is 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 auditing for processes that start prior to auditd is enabled	☐	☐
6.2.1.3	Ensure audit_backlog_limit is configured	☐	☐
6.2.1.4	Ensure auditd service is enabled and active	☐	☐
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	☐	☐
6.2.3.1	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 /etc/sysconfig/network and /etc/NetworkManager/system-connections/ 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	☐	☐

Recommendation		Set Correctly	
		Yes	No
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 are collected	☐	☐
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	☐	☐

Recommendation		Set Correctly	
		Yes	No
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 the audit log file directory mode is configured	☐	☐
6.2.4.2	Ensure audit log files mode is configured	☐	☐
6.2.4.3	Ensure audit log files owner is configured	☐	☐
6.2.4.4	Ensure audit log files group owner 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	☐	☐
6.3.3	Ensure cryptographic mechanisms are used to protect the integrity of audit tools	☐	☐
6.3.4	Ensure report_url 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	☐	☐

Recommendation		Set Correctly	
		Yes	No
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	☐	☐
7.2.8	Ensure local interactive user home directories are configured	☐	☐
7.2.9	Ensure local interactive user dot files access is configured	☐	☐

Appendix: CIS Controls v8 Unmapped Recommendations

Recommendation		Set Correctly	
		Yes	No
1.1.2.5.1	Ensure separate partition exists for /var/tmp	☐	☐
1.4.3	Ensure authentication is required when booting into rescue mode	☐	☐
1.4.4	Ensure authentication is required when booting into emergency mode	☐	☐
1.4.5	Ensure systemd ability to spawn an interactive boot process is disabled	☐	☐
1.5.5	Ensure kernel.kptr_restrict is configured	☐	☐
3.3.1.20	Ensure net.ipv4.conf.default.route_localnet is configured	☐	☐
4.1.8	Ensure firewalld default zone is set to drop	☐	☐
4.1.9	Ensure conflicting firewall utilities are not in use with firewalld	☐	☐
4.1.10	Ensure firewalld logs denied packets	☐	☐
4.1.11	Ensure firewalld blocks ICMP timestamp requests and replies	☐	☐
5.1.2	Ensure access to /etc/ssh/sshd_config.d/ directory is configured	☐	☐
6.1.2.10	Ensure a rsyslog TLS network stream driver is installed	☐	☐
6.1.2.11	Ensure rsyslog network forwarding uses TLS	☐	☐
6.1.2.12	Ensure rsyslog TLS CA certificate is configured	☐	☐

Appendix: Change History

Date	Version	Changes for this version
2026-08-11	3.0.0	**Added (92):**
		1.1.1.6 Ensure overlay kernel module is not available
		1.1.1.9 Ensure firewire-core kernel module is not available
		1.2.1.5 Ensure weak dependencies are configured
		1.4.3 Ensure authentication is required when booting into rescue mode
		1.4.4 Ensure authentication is required when booting into emergency mode
		1.4.5 Ensure systemd ability to spawn an interactive boot process is disabled
		1.5.1 Ensure fs.protected_hardlinks is configured
		1.5.2 Ensure fs.protected_symlinks is configured
		1.5.3 Ensure fs.suid_dumpable is configured
		1.5.4 Ensure kernel.dmesg_restrict is configured
		1.5.5 Ensure kernel.kptr_restrict is configured
		1.5.8 Ensure core file size is configured
		1.5.9 Ensure systemd-coredump ProcessSizeMax is configured
		1.6.7 Ensure system wide crypto policy enables post-quantum algorithms
		1.8.7 Ensure Wayland is configured
		2.1.3 Ensure cockpit web services are not in use
		2.1.22 Ensure Xwayland is not in use
		2.4.1.7 Ensure access to /etc/cron.yearly is configured

Date	Version	Changes for this version
		3.2.5 Ensure atm kernel module is not available
		3.2.6 Ensure can kernel module is 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
		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

Date	Version	Changes for this version
		3.3.1.18 Ensure net.ipv4.tcp_syncookies is configured
		3.3.1.19 Ensure net.ipv4.conf.all.route_localnet is configured
		3.3.1.20 Ensure net.ipv4.conf.default.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 firewalld is installed
		4.1.2 Ensure firewalld backend is configured
		4.1.3 Ensure firewalld.service is configured
		4.1.4 Ensure firewalld active zone target is configured
		4.1.6 Ensure firewalld loopback source address traffic is configured
		4.1.8 Ensure firewalld default zone is set to drop
		4.1.9 Ensure conflicting firewall utilities are not in use with firewalld
		4.1.10 Ensure firewalld logs denied packets

Date	Version	Changes for this version
		4.1.11 Ensure firewalld blocks ICMP timestamp requests and replies
		5.1.2 Ensure access to /etc/ssh/sshd_config.d/ directory is configured
		5.2.2 Ensure access on sudoers configuration files is configured
		5.2.3 Ensure access on the /etc/sudoers.d directory is configured
		5.2.8 Ensure sudo pwfeedback is configured
		5.2.9 Ensure sudo visiblepw is configured
		5.2.11 Ensure sudo timestamp_type is configured
		5.4.3.3 Ensure default user umask is configured in /etc/login.defs
		6.1.2.4 Ensure access to rsyslog configuration files is configured
		6.1.2.5 Ensure access to the /etc/rsyslog.d directory is configured
		6.1.2.6 Ensure access to the rsyslog work directory is configured
		6.1.2.10 Ensure a rsyslog TLS network stream driver is installed
		6.1.2.11 Ensure rsyslog network forwarding uses TLS
		6.1.2.12 Ensure rsyslog TLS CA certificate is configured
		6.1.2.13 Ensure rsyslog TLS client certificate is configured
		6.2.3.1 Ensure modification of the /etc/sudoers file is collected
		6.2.3.5 Ensure events that modify sethostname and setdomainname are collected

Date	Version	Changes for this version
		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 /etc/sysconfig/network and /etc/NetworkManager/system-connections/ are collected
		6.2.3.9 Ensure events that modify /etc/NetworkManager directory 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 are collected
		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
		6.2.3.24 Ensure unlink file deletion events by users are collected

Date	Version	Changes for this version
		6.2.3.25 Ensure rename file deletion events by users are 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
		6.3.4 Ensure report_url is configured
		Removed / Restructured (34):
		1.5.3 Ensure core dump backtraces are disabled restructured into the systemd-coredump recommendations 1.5.8-1.5.10 (core file size / ProcessSizeMax / Storage)
		1.6.2 Ensure system wide crypto policy is not set in sshd configuration removed - SSH cipher/MAC/KexAlgorithms selection now defers to the system-wide crypto policy (section 1.6), making a separate not-set-in-sshd recommendation redundant
		2.1.10 Ensure nis server services are not in use removed - the NIS server package (ypserv) was removed from EL9 by the vendor and is not available on any covered platform; auditing an uninstallable package has no security merit
		2.1.19 Ensure xinetd services are not in use removed - xinetd is not shipped on EL9; the recommendation audited a package that cannot be installed on the platform
		2.2.3 Ensure nis client is not installed removed - the NIS client package (ypbind) was removed from EL9 and is not available on any covered platform (including Amazon Linux 2023); auditing an uninstallable package has no security merit

Date	Version	Changes for this version
		3.3.1 Ensure ip forwarding is disabled restructured into per-parameter recommendations 3.3.1.1-.3 (net.ipv4.ip_forward / conf.all|default.forwarding) + 3.3.2.1-.2 (IPv6)
		3.3.2 Ensure packet redirect sending is disabled restructured into 3.3.1.4/.5 (net.ipv4.conf.all|default.send_redirects)
		3.3.3 Ensure bogus icmp responses are ignored restructured into 3.3.1.6 (net.ipv4.icmp_ignore_bogus_error_responses)
		3.3.4 Ensure broadcast icmp requests are ignored restructured into 3.3.1.7 (net.ipv4.icmp_echo_ignore_broadcasts)
		3.3.5 Ensure icmp redirects are not accepted restructured into 3.3.1.8/.9 (net.ipv4 accept_redirects) + 3.3.2.3/.4 (IPv6)
		3.3.6 Ensure secure icmp redirects are not accepted restructured into 3.3.1.10/.11 (net.ipv4.conf.all|default.secure_redirects)
		3.3.7 Ensure reverse path filtering is enabled restructured into 3.3.1.12/.13 (net.ipv4.conf.all|default.rp_filter)
		3.3.8 Ensure source routed packets are not accepted restructured into 3.3.1.14/.15 (net.ipv4 accept_source_route) + 3.3.2.5/.6 (IPv6)
		3.3.9 Ensure suspicious packets are logged restructured into 3.3.1.16/.17 (net.ipv4.conf.all|default.log_martians)
		3.3.10 Ensure tcp syn cookies is enabled restructured into 3.3.1.18 (net.ipv4.tcp_syncookies)

Date	Version	Changes for this version
		3.3.11 Ensure ipv6 router advertisements are not accepted restructured into 3.3.2.7/.8 (net.ipv6.conf.all|default.accept_ra)
		4.1.1 Ensure nftables is installed removed - the benchmark consolidated on firewalld as the single host firewall (firewalld manages nftables as its backend, see 4.1.2); the standalone nftables recommendations were dropped
		4.1.2 Ensure a single firewall configuration utility is in use removed - superseded by the firewalld-focused rework; the single-firewall-utility intent is now covered by 4.1.9 (conflicting firewall utilities are not in use)
		4.3.1 Ensure nftables base chains exist removed - dropped with the standalone nftables path; the host firewall is now configured through firewalld (section 4.1)
		4.3.2 Ensure nftables established connections are configured removed - dropped with the standalone nftables path (see section 4.1)
		4.3.3 Ensure nftables default deny firewall policy removed - dropped with the standalone nftables path (see section 4.1)
		4.3.4 Ensure nftables loopback traffic is configured removed - dropped with the standalone nftables path (see section 4.1)
		6.2.1.4 Ensure only one logging system is in use removed - superseded by the journald + rsyslog coexistence model: v3.0.0 runs both (journald forwarding to rsyslog) rather than requiring a single logging system
		6.2.2.1.1 Ensure systemd-journal-remote is installed removed - the systemd-journal-remote/upload remote-logging path was dropped in favor of rsyslog TLS forwarding (section 6.1.2) as the remote-log mechanism

Date	Version	Changes for this version
		6.2.2.1.2 Ensure systemd-journal-upload authentication is configured removed - dropped with the systemd-journal-upload path (remote logging is via rsyslog TLS forwarding, section 6.1.2)
		6.2.2.1.3 Ensure systemd-journal-upload is enabled and active removed - dropped with the systemd-journal-upload path (remote logging is via rsyslog TLS forwarding, section 6.1.2)
		6.2.2.2 Ensure journald ForwardToSyslog is disabled removed - dropped in the logging restructure; it enforced the keep-journald-standalone model. v3.0.0 forwards journald to rsyslog instead (new 6.1.1.3)
		6.2.2.3 Ensure journald Compress is configured removed - existed to support a journald-only logging configuration; unnecessary under the v3.0.0 restructure to rsyslog (journald now forwards to rsyslog, which provides persistent storage)
		6.2.2.4 Ensure journald Storage is configured removed - existed to support a journald-only logging configuration; unnecessary under the v3.0.0 restructure to rsyslog (journald now forwards to rsyslog, which provides persistent storage)
		6.3.3.1 Ensure changes to system administration scope (sudoers) is collected restructured into 6.2.3.1 (modification of /etc/sudoers is collected)
		6.3.3.5 Ensure events that modify the system's network environment are collected restructured/split into 6.2.3.5-.9 (sethostname, /etc/issue, /etc/hosts, sysconfig/network, NetworkManager)
		6.3.3.8 Ensure events that modify user/group information are collected restructured/split into 6.2.3.12-.17 (/etc/group, /etc/passwd, /etc/shadow, opasswd, nsswitch, pam)

Date	Version	Changes for this version
		6.3.3.9 Ensure discretionary access control permission modification events are collected restructured/split into 6.2.3.18-.20 (chmod / chown / setxattr families)
		6.3.3.13 Ensure file deletion events by users are collected restructured/split into 6.2.3.24/.25 (unlink / rename deletion events)
		Retitled (51):
		1.1.2.1.1: "Ensure /tmp is a separate partition" -> "Ensure /tmp is tmpfs or a separate partition"
		1.1.2.2.1: "Ensure /dev/shm is a separate partition" -> "Ensure /dev/shm is tmpfs or a separate partition"
		1.2.1.2: "Ensure gpgcheck is globally activated" -> "Ensure gpgcheck is configured"
		1.2.1.3: "Ensure repo_gpgcheck is globally activated" -> "Ensure repo_gpgcheck is configured"
		1.5.6: "Ensure address space layout randomization is enabled" -> "Ensure kernel.randomize_va_space is configured"
		1.5.7: "Ensure ptrace_scope is restricted" -> "Ensure kernel.yama.ptrace_scope is configured"
		1.5.10: "Ensure core dump storage is disabled" -> "Ensure systemd-coredump Storage is configured"
		1.6.1: "Ensure system wide crypto policy is not set to legacy" -> "Ensure system wide crypto policy is not legacy"
		1.6.3: "Ensure system wide crypto policy disables macs less than 128 bits" -> "Ensure system wide crypto policy macs are configured"
		1.7.1: "Ensure message of the day is configured properly" -> "Ensure /etc/motd is configured"
		1.7.2: "Ensure local login warning banner is configured properly" -> "Ensure /etc/issue is configured"

Date	Version	Changes for this version
		1.7.3: "Ensure remote login warning banner is configured properly" -> "Ensure /etc/issue.net is configured"
		1.8.2: "Ensure GDM disable-user-list option is enabled" -> "Ensure GDM disable-user-list is configured"
		1.8.3: "Ensure GDM screen locks when the user is idle" -> "Ensure GDM screen lock is configured"
		1.8.3: "Ensure GDM screen locks cannot be overridden" -> "Ensure GDM screen lock is configured"
		1.8.4: "Ensure GDM automatic mounting of removable media is disabled" -> "Ensure GDM automount is configured"
		1.8.4: "Ensure GDM disabling automatic mounting of removable media is not overridden" -> "Ensure GDM automount is configured"
		1.8.5: "Ensure GDM autorun-never is enabled" -> "Ensure GDM autorun-never is configured"
		1.8.5: "Ensure GDM autorun-never is not overridden" -> "Ensure GDM autorun-never is configured"
		2.1.20: "Ensure GNOME Display Manager is removed" -> "Ensure a graphical user interface is not in use"
		2.4.1.2: "Ensure permissions on /etc/crontab are configured" -> "Ensure access to /etc/crontab is configured"
		2.4.1.3: "Ensure permissions on /etc/cron.hourly are configured" -> "Ensure access to /etc/cron.hourly is configured"
		2.4.1.4: "Ensure permissions on /etc/cron.daily are configured" -> "Ensure access to /etc/cron.daily is configured"
		2.4.1.5: "Ensure permissions on /etc/cron.weekly are configured" -> "Ensure access to /etc/cron.weekly is configured"
		2.4.1.6: "Ensure permissions on /etc/cron.monthly are configured" -> "Ensure access to /etc/cron.monthly is configured"

Date	Version	Changes for this version
		2.4.1.8: "Ensure permissions on /etc/cron.d are configured" -> "Ensure access to /etc/cron.d is configured"
		2.4.1.9: "Ensure crontab is restricted to authorized users" -> "Ensure access to crontab is configured"
		2.4.2.1: "Ensure at is restricted to authorized users" -> "Ensure access to at is configured"
		3.1.2: "Ensure wireless interfaces are disabled" -> "Ensure wireless interfaces are not available"
		4.1.7: "Ensure firewalld drops unnecessary services and ports" -> "Ensure firewalld services and ports are configured"
		5.1.1: "Ensure permissions on /etc/ssh/sshd_config are configured" -> "Ensure access to /etc/ssh/sshd_config is configured"
		5.1.3: "Ensure permissions on SSH private host key files are configured" -> "Ensure access to SSH private host key files is configured"
		5.1.4: "Ensure permissions on SSH public host key files are configured" -> "Ensure access to SSH public host key files is configured"
		5.2.10: "Ensure sudo authentication timeout is configured correctly" -> "Ensure sudo timestamp_timeout is configured"
		5.3.2.1: "Ensure active authselect profile includes pam modules" -> "Ensure PAM is managed by an active authselect profile"
		6.1.1.1: "Ensure journald service is enabled and active" -> "Ensure journald service is active"
		6.1.1.2: "Ensure systemd-journal-remote service is not in use" -> "Ensure systemd-journal-remote is not in use"
		6.1.2.7: "Ensure rsyslog is not configured to receive logs from a remote client" -> "Ensure rsyslog is not configured as a log server"

Date	Version	Changes for this version
		6.1.2.8: "Ensure rsyslog logging is configured" -> "Ensure rsyslog logging rules are configured"
		6.1.2.14: "Ensure rsyslog logrotate is configured" -> "Ensure logrotate is configured"
		6.2.1.3: "Ensure audit_backlog_limit is sufficient" -> "Ensure audit_backlog_limit is configured"
		7.1.1: "Ensure permissions on /etc/passwd are configured" -> "Ensure access to /etc/passwd is configured"
		7.1.2: "Ensure permissions on /etc/passwd- are configured" -> "Ensure access to /etc/passwd- is configured"
		7.1.3: "Ensure permissions on /etc/group are configured" -> "Ensure access to /etc/group is configured"
		7.1.4: "Ensure permissions on /etc/group- are configured" -> "Ensure access to /etc/group- is configured"
		7.1.5: "Ensure permissions on /etc/shadow are configured" -> "Ensure access to /etc/shadow is configured"
		7.1.6: "Ensure permissions on /etc/shadow- are configured" -> "Ensure access to /etc/shadow- is configured"
		7.1.7: "Ensure permissions on /etc/gshadow are configured" -> "Ensure access to /etc/gshadow is configured"
		7.1.8: "Ensure permissions on /etc/gshadow- are configured" -> "Ensure access to /etc/gshadow- is configured"
		7.1.9: "Ensure permissions on /etc/shells are configured" -> "Ensure access to /etc/shells is configured"
		7.1.10: "Ensure permissions on /etc/security/opasswd are configured" -> "Ensure access to /etc/security/opasswd is configured"
		Renumbered (116):
		1.1.1.6 -> 1.1.1.7 Ensure squashfs kernel module is not available

Date	Version	Changes for this version
		1.1.1.7 -> 1.1.1.8 Ensure udf kernel module is not available
		1.1.1.8 -> 1.1.1.10 Ensure usb-storage kernel module is not available
		1.1.1.9 -> 1.1.1.11 Ensure unused filesystems kernel modules are not available
		1.5.1 -> 1.5.6 Ensure kernel.randomize_va_space is configured
		1.5.2 -> 1.5.7 Ensure kernel.yama.ptrace_scope is configured
		1.5.4 -> 1.5.10 Ensure systemd-coredump Storage is configured
		1.6.3 -> 1.6.2 Ensure system wide crypto policy disables sha1 hash and signature support
		1.6.4 -> 1.6.3 Ensure system wide crypto policy macs are configured
		1.6.5 -> 1.6.4 Ensure system wide crypto policy disables cbc for ssh
		1.6.6 -> 1.6.5 Ensure system wide crypto policy disables chacha20-poly1305 for ssh
		1.6.7 -> 1.6.6 Ensure system wide crypto policy disables EtM for ssh
		1.8.2 -> 1.8.1 Ensure GDM login banner is configured
		1.8.3 -> 1.8.2 Ensure GDM disable-user-list is configured
		1.8.4 -> 1.8.3 Ensure GDM screen lock is configured
		1.8.5 -> 1.8.3 Ensure GDM screen lock is configured
		1.8.6 -> 1.8.4 Ensure GDM automount is configured
		1.8.7 -> 1.8.4 Ensure GDM automount is configured
		1.8.8 -> 1.8.5 Ensure GDM autorun-never is configured

Date	Version	Changes for this version
		1.8.9 -> 1.8.5 Ensure GDM autorun-never is configured
		1.8.10 -> 1.8.6 Ensure XDMCP is not enabled
		2.1.3 -> 2.1.4 Ensure dhcp server services are not in use
		2.1.4 -> 2.1.5 Ensure dns server services are not in use
		2.1.5 -> 2.1.6 Ensure dnsmasq services are not in use
		2.1.6 -> 2.1.7 Ensure samba file server services are not in use
		2.1.7 -> 2.1.8 Ensure ftp server services are not in use
		2.1.8 -> 2.1.9 Ensure message access server services are not in use
		2.1.9 -> 2.1.10 Ensure network file system services are not in use
		2.1.21 -> 2.1.19 Ensure mail transfer agents are configured for local-only mode
		1.8.1 -> 2.1.20 Ensure a graphical user interface is not in use
		2.1.20 -> 2.1.21 Ensure X window server services are not in use
		2.1.22 -> 2.1.23 Ensure only approved services are listening on a network interface
		2.2.4 -> 2.2.3 Ensure telnet client is not installed
		2.2.5 -> 2.2.4 Ensure tftp client is not installed
		2.4.1.7 -> 2.4.1.8 Ensure access to /etc/cron.d is configured
		2.4.1.8 -> 2.4.1.9 Ensure access to crontab is configured
		4.2.2 -> 4.1.5 Ensure firewalld loopback traffic is configured
		4.2.1 -> 4.1.7 Ensure firewalld services and ports are configured

Date	Version	Changes for this version
		5.1.2 -> 5.1.3 Ensure access to SSH private host key files is configured
		5.1.3 -> 5.1.4 Ensure access to SSH public host key files is configured
		5.1.4 -> 5.1.5 Ensure sshd Ciphers are configured
		5.1.5 -> 5.1.6 Ensure sshd KexAlgorithms is configured
		5.1.6 -> 5.1.7 Ensure sshd MACs are configured
		5.1.7 -> 5.1.8 Ensure sshd access is configured
		5.1.8 -> 5.1.9 Ensure sshd Banner is configured
		5.1.9 -> 5.1.10 Ensure sshd ClientAliveInterval and ClientAliveCountMax are configured
		5.1.10 -> 5.1.11 Ensure sshd DisableForwarding is enabled
		5.1.11 -> 5.1.12 Ensure sshd GSSAPIAuthentication is disabled
		5.1.12 -> 5.1.13 Ensure sshd HostbasedAuthentication is disabled
		5.1.13 -> 5.1.14 Ensure sshd IgnoreRhosts is enabled
		5.1.14 -> 5.1.15 Ensure sshd LoginGraceTime is configured
		5.1.15 -> 5.1.16 Ensure sshd LogLevel is configured
		5.1.16 -> 5.1.17 Ensure sshd MaxAuthTries is configured
		5.1.17 -> 5.1.18 Ensure sshd MaxStartups is configured
		5.1.18 -> 5.1.19 Ensure sshd MaxSessions is configured
		5.1.19 -> 5.1.20 Ensure sshd PermitEmptyPasswords is disabled
		5.1.20 -> 5.1.21 Ensure sshd PermitRootLogin is disabled

Date	Version	Changes for this version
		5.1.21 -> 5.1.22 Ensure sshd PermitUserEnvironment is disabled
		5.1.22 -> 5.1.23 Ensure sshd UsePAM is enabled
		5.2.2 -> 5.2.4 Ensure sudo commands use pty
		5.2.3 -> 5.2.5 Ensure sudo log file exists
		5.2.4 -> 5.2.6 Ensure users must provide password for escalation
		5.2.5 -> 5.2.7 Ensure re-authentication for privilege escalation is not disabled globally
		5.2.6 -> 5.2.10 Ensure sudo timestamp_timeout is configured
		5.2.7 -> 5.2.12 Ensure access to the su command is restricted
		5.4.3.3 -> 5.4.3.4 Ensure default user umask is configured
		6.2.1.1 -> 6.1.1.1 Ensure journald service is active
		6.2.2.1.4 -> 6.1.1.2 Ensure systemd-journal-remote is not in use
		6.2.3.3 -> 6.1.1.3 Ensure journald is configured to send logs to rsyslog
		6.2.1.3 -> 6.1.1.4 Ensure journald log file rotation is configured
		6.2.1.2 -> 6.1.1.5 Ensure journald log file access is configured
		6.2.3.1 -> 6.1.2.1 Ensure rsyslog is installed
		6.2.3.2 -> 6.1.2.2 Ensure rsyslog service is enabled and active
		6.2.3.4 -> 6.1.2.3 Ensure rsyslog log file creation mode is configured
		6.2.3.7 -> 6.1.2.7 Ensure rsyslog is not configured as a log server
		6.2.3.5 -> 6.1.2.8 Ensure rsyslog logging rules are configured

Date	Version	Changes for this version
		6.2.3.6 -> 6.1.2.9 Ensure rsyslog is configured to send logs to a remote log host
		6.2.3.8 -> 6.1.2.14 Ensure logrotate is configured
		6.2.4.1 -> 6.1.3.1 Ensure access to all logfiles has been configured
		6.3.1.1 -> 6.2.1.1 Ensure auditd packages are installed
		6.3.1.2 -> 6.2.1.2 Ensure auditing for processes that start prior to auditd is enabled
		6.3.1.3 -> 6.2.1.3 Ensure audit_backlog_limit is configured
		6.3.1.4 -> 6.2.1.4 Ensure auditd service is enabled and active
		6.3.2.1 -> 6.2.2.1 Ensure audit log storage size is configured
		6.3.2.2 -> 6.2.2.2 Ensure audit logs are not automatically deleted
		6.3.2.3 -> 6.2.2.3 Ensure system is disabled when audit logs are full
		6.3.2.4 -> 6.2.2.4 Ensure system warns when audit logs are low on space
		6.3.3.2 -> 6.2.3.2 Ensure actions as another user are always logged
		6.3.3.3 -> 6.2.3.3 Ensure events that modify the sudo log file are collected
		6.3.3.4 -> 6.2.3.4 Ensure events that modify date and time information are collected
		6.3.3.6 -> 6.2.3.10 Ensure use of privileged commands are collected
		6.3.3.7 -> 6.2.3.11 Ensure unsuccessful file access attempts are collected
		6.3.3.10 -> 6.2.3.21 Ensure successful file system mounts are collected

Date	Version	Changes for this version
		6.3.3.11 -> 6.2.3.22 Ensure session initiation information is collected
		6.3.3.12 -> 6.2.3.23 Ensure login and logout events are collected
		6.3.3.14 -> 6.2.3.26 Ensure events that modify the system's Mandatory Access Controls are collected
		6.3.3.15 -> 6.2.3.27 Ensure successful and unsuccessful attempts to use the chcon command are collected
		6.3.3.16 -> 6.2.3.28 Ensure successful and unsuccessful attempts to use the setfacl command are collected
		6.3.3.17 -> 6.2.3.29 Ensure successful and unsuccessful attempts to use the chacl command are collected
		6.3.3.18 -> 6.2.3.30 Ensure successful and unsuccessful attempts to use the usermod command are collected
		6.3.3.19 -> 6.2.3.31 Ensure kernel module loading unloading and modification is collected
		6.3.3.20 -> 6.2.3.35 Ensure the audit configuration is immutable
		6.3.3.21 -> 6.2.3.36 Ensure the running and on disk configuration is the same
		6.3.4.1 -> 6.2.4.1 Ensure the audit log file directory mode is configured
		6.3.4.2 -> 6.2.4.2 Ensure audit log files mode is configured
		6.3.4.3 -> 6.2.4.3 Ensure audit log files owner is configured
		6.3.4.4 -> 6.2.4.4 Ensure audit log files group owner is configured
		6.3.4.5 -> 6.2.4.5 Ensure audit configuration files mode is configured
		6.3.4.6 -> 6.2.4.6 Ensure audit configuration files owner is configured

Date	Version	Changes for this version
		6.3.4.7 -> 6.2.4.7 Ensure audit configuration files group owner is configured
		6.3.4.8 -> 6.2.4.8 Ensure audit tools mode is configured
		6.3.4.9 -> 6.2.4.9 Ensure audit tools owner is configured
		6.3.4.10 -> 6.2.4.10 Ensure audit tools group owner is configured
		6.1.1 -> 6.3.1 Ensure AIDE is installed
		6.1.2 -> 6.3.2 Ensure filesystem integrity is regularly checked
		6.1.3 -> 6.3.3 Ensure cryptographic mechanisms are used to protect the integrity of audit tools