INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-05-30 19:07 UTC

CIFSwitch: 19-Year-Old Linux Kernel CIFS Flaw Enables Local Privilege Escalation to Root

CVE VULNERABILITY | HIGH | CVSS 7.5

SCC Item ID	SCC-CVE-2026-0242
Type	CVE Vulnerability
Severity	HIGH
CVSS Base Score	7.5
Affected Products	Linux kernel (CIFS subsystem) with cifs-utils 6.14+; confirmed: Linux Mint 21.3/22.3, CentOS Stream 9, Rocky Linux 9, AlmaLinux 9, Kali Linux 2021.4-2026.1, SLES 15 SP7; potentially affected: Ubuntu, Debian, Pop!_OS, openSUSE, Oracle Linux, Amazon Linux (with cifs-utils installed)
Published	2026-05-30T10:16:08
Discovery Source	Rss

Executive Summary

A 19-year-old privilege escalation flaw in the Linux kernel's CIFS subsystem, dubbed CIFSwitch, allows any unprivileged local user to gain full root access on affected systems. Enterprise Linux distributions including CentOS Stream 9, Rocky Linux 9, AlmaLinux 9, and SLES 15 SP7 are confirmed vulnerable in default configurations, and a public proof-of-concept exploit is already available. (Note: As of 2026-05-28, no CVE identifier had been formally assigned; a CVE may be assigned retroactively. Monitor NVD and vendor advisories for assignment.) Organizations running these distributions with cifs-utils installed face immediate risk of complete system compromise from any authenticated local user or process.

Technical Analysis

CIFSwitch is a local privilege escalation vulnerability residing in the Linux kernel's CIFS subsystem, present since the 2007 introduction of that code path. No CVE identifier had been assigned as of the reporting date (2026-05-28). Relevant CWEs: CWE-269 (Improper Privilege Management), CWE-346 (Origin Validation Error), CWE-284 (Improper Access Control). MITRE ATT&CK techniques: T1068 (Exploitation for Privilege Escalation), T1055 (Process Injection), T1574.006 (Hijack Execution Flow: Dynamic Linker Hijacking), T1548.001 (Abuse Elevation Control Mechanism: Setuid and Setgid). Attack vector: an unprivileged local user forges Kerberos/SPNEGO key requests to the kernel, triggering the kernel's request-key mechanism and causing the cifs-utils userspace helper to load an attacker-supplied NSS (Name Service Switch) module. Because this

helper runs with elevated privileges, loading an attacker-controlled shared library results in arbitrary code execution as root. Exploitation requires user namespaces to be enabled, which is the default on CentOS Stream 9, Rocky Linux 9, AlmaLinux 9, Ubuntu 20.04+, and Debian 11+. Check kernel parameters 'kernel.unprivileged_userns_clone' or 'user.max_user_namespaces' on your systems to verify status. Exploitation also requires cifs-utils installed. Confirmed affected: Linux Mint 21.3/22.3, CentOS Stream 9, Rocky Linux 9, AlmaLinux 9, Kali Linux 2021.4-2026.1, SLES 15 SP7 with cifs-utils 6.14+. Potentially affected: Ubuntu, Debian, Pop!_OS, openSUSE, Oracle Linux, Amazon Linux where cifs-utils is installed. A public proof-of-concept exploit is available, significantly lowering exploitation complexity. AlmaLinux published a call for patched kernel testing on 2026-05-28; no upstream kernel patch or CVE was confirmed assigned at time of reporting. CVSS base score: 7.5 (High).

Action Checklist

1. Step 1: Containment, Immediately inventory all Linux systems running cifs-utils 6.14+ across CentOS Stream 9, Rocky Linux 9, AlmaLinux 9, SLES 15 SP7, and any Debian/Ubuntu-family hosts where cifs-utils is installed. Where cifs-utils is not operationally required, remove or disable it (e.g., 'yum remove cifs-utils' or 'apt remove cifs-utils') until a vendor patch is available. Where removal is not feasible, restrict local interactive user access to affected hosts and limit who can authenticate locally (NIST AC-6, Least Privilege; CIS 5.4, Restrict Administrator Privileges).
2. Step 2: Detection, First, verify auditd is installed and running: systemctl status auditd. If not installed, install via yum install audit (RHEL-family) or apt install auditd (Debian-family) and ensure the service is enabled. Query endpoint and SIEM telemetry for unexpected loading of NSS modules by kernel helper processes (look for '/sbin/request-key' or 'cifs.upcall' spawning unusual child processes or loading shared libraries from non-standard paths). Monitor audit logs for privilege escalation events: ausearch -m USER_AUTH,PRIV_ESC,SYSCALL -ts recent. Enable auditd rules tracking execve and mmap syscalls by the cifs.upcall helper. Check for unusual entries in /etc/nsswitch.conf or unexpected .so files in NSS library paths. Reference: NIST AU-2 (Event Logging), NIST SI-4 (System Monitoring), CIS 8.2 (Collect Audit Logs).
3. Step 3: Eradication, Apply vendor-issued patched kernels as they become available. AlmaLinux has published patched kernels for community testing as of 2026-05-28 (see https://almalinux.org/blog/2026-05-28-cifswitch/); monitor CentOS Stream, Rocky Linux, SUSE, and upstream kernel channels for corresponding fixes. If cifs-utils is not required for operations, remove it permanently. If CIFS mounts are required, restrict user namespace creation where feasible: set 'kernel.unprivileged_userns_clone=0' (Debian-family) or 'user.max_user_namespaces=0' (RHEL-family) via sysctl as a temporary mitigation, validate operational impact before deploying. Reference: NIST SI-2 (Flaw Remediation), CIS 7.3 (Perform Automated OS Patch Management).
4. Step 4: Recovery, After patching, verify the installed kernel version matches the vendor-issued fix and that cifs-utils has been updated or removed as appropriate. Re-enable any services paused during containment. Run a post-remediation audit: confirm no unauthorized accounts were created, no cron jobs or startup scripts were modified, and no new SUID binaries exist (find / -perm -4000 -type f 2>/dev/null). Review auditd logs for any exploitation activity that occurred before patching. Reference: NIST AU-6 (Audit Record Review), NIST IR-4 (Incident Handling), NIST CA-7 (Continuous Monitoring).
5. Step 5: Post-Incident, Conduct a control gap review: assess whether least-privilege policies prevented unnecessary local user access to affected hosts (NIST AC-6), whether software inventory practices would have flagged cifs-utils as an unnecessary installed package (CIS 2.1, Establish and Maintain a Software

Inventory), and whether user namespace restrictions are standardized in hardening baselines (CIS 4.6, Securely Manage Enterprise Assets and Software). Update vulnerability management SLAs to account for unassigned-CVE threats with public PoC exploits. Reference: NIST CM-6 (Configuration Settings), CIS 7.1 (Establish and Maintain a Vulnerability Management Process).

IR / Forensic Enrichment

Triage Priority	IMMEDIATE
Escalation Criteria	Escalate to CISO and initiate full incident response if auditd or process monitoring detects execve events spawned from cifs.upcall or request-key with a root-owned resulting process, any new UID 0 account appears in /etc/passwd, SUID binaries are found with modification timestamps during the exposure window, or if affected hosts process PII, PHI, or payment card data — triggering potential breach notification obligations under HIPAA, GDPR, or PCI-DSS.
Recovery Notes	After applying vendor-patched kernels, verify each host with 'uname -r' against the specific kernel version cited in the distribution advisory (AlmaLinux 2026-05-28 advisory as initial reference; track CentOS Stream, Rocky Linux, and SUSE channels for their corresponding fixes). Maintain the user.max_user_namespaces=0 sysctl even after kernel patching until all hosts in scope are confirmed patched, then reassess whether the namespace restriction should remain as a defense-in-depth hardening control. Monitor auditd logs for at least 30 days post-patching for deferred exploitation indicators — specifically any anomalous root session activity or NSS library file modifications that would indicate a threat actor with prior access attempting to activate a pre-planted payload.
Forensic Artifacts	/etc/nsswitch.conf modification timestamps and content — CIFSwitch exploits the NSS module loading path invoked by cifs.upcall; an attacker staging the exploit may modify nsswitch.conf to insert a malicious resolver or place a rogue libnss_*.so in a path that takes precedence over legitimate NSS libraries | Auditd SYSCALL records for execve and mmap syscalls with exe=/sbin/cifs.upcall or exe=/sbin/request-key as the audited process — the exploit mechanism requires the kernel to invoke cifs.upcall in a context where attacker-controlled NSS resolution leads to execution of attacker code, making these the primary forensic indicator of exploitation | World-writable directory contents in /tmp, /var/tmp, and /dev/shm for *.so files, compiled ELF binaries, or scripts — the public PoC for CIFSwitch stages exploit components in these directories prior to triggering the cifs.upcall NSS loading path, and artifact timestamps here will correlate with the exploitation attempt | SUID binary inventory delta — a successful CIFSwitch exploitation grants full root; attackers commonly install a SUID shell (/bin/bash -p wrapper or a copy of /bin/sh with SUID set) as a persistence backdoor immediately after gaining root, detectable via 'find / -perm -4000 -type f' compared against a known-good baseline | /var/log/secure or /var/log/auth.log entries showing privilege changes (su, sudo, newuidmap, newgidmap) from non-administrative accounts during the exposure window, correlated with cifs.upcall invocation timestamps from journald or auditd — these log entries establish the forensic chain linking an unprivileged user account to the root escalation event

Per-Action IR Details

Step 1: Containment — Immediately inventory all Linux systems running cifs-utils 6.14+ across CentOS Stream 9, Rocky Linux 9, AlmaLinux 9, SLES 15 SP7, and any Debian/Ubuntu-family hosts where cifs-utils is installed. Where cifs-utils is not operationally required, remove or disable it (e.g., 'yum remove cifs-utils' or 'apt remove cifs-utils') until a vendor patch is available. Where removal is not feasible, restrict local interactive

user access to affected hosts and limit who can authenticate locally (NIST AC-6 — Least Privilege; CIS 5.4 — Restrict Administrator Privileges).

NIST Phase: Containment

Reference: NIST 800-61r3 §3.3 — Containment Strategy

Controls: NIST AC-6 (Least Privilege), NIST CM-7 (Least Functionality), CIS 5.4 (Restrict Administrator Privileges to Dedicated Administrator Accounts), CIS 2.3 (Address Unauthorized Software), CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory)

Compensating: Run the following one-liner across all managed Linux hosts via SSH or Ansible ad-hoc to identify exposed systems: 'rpm -qa cifs-utils 2>/dev/null || dpkg -l cifs-utils 2>/dev/null | grep ^ii'. For systems where cifs-utils cannot be removed, immediately restrict console and SSH access to named admin accounts only via /etc/security/access.conf or PAM: add '-:ALL EXCEPT root wheel:LOCAL' to /etc/security/access.conf and restart the PAM-controlled services. Use 'getent passwd | awk -F: "$3 >= 1000"' to enumerate all interactive user accounts that could trigger the exploit on each host.

Evidence: Before modifying any system, capture: (1) 'rpm -qa --last cifs-utils' or 'dpkg -l cifs-utils' to document installed version and install date; (2) 'cat /etc/nsswitch.conf' to record current NSS resolver order — exploitation of CIFSwitch involves NSS module loading via cifs.upcall, so the pre-containment nsswitch.conf state is a key forensic baseline; (3) 'ls -la /lib/x86_64-linux-gnu/libnss_*.so* /usr/lib64/libnss_*.so*' to document all currently registered NSS shared libraries and detect any injected .so files placed by an attacker prior to your response; (4) 'find /tmp /var/tmp /dev/shm -name "*.so" -o -name "request-key*" 2>/dev/null' to identify exploit staging artifacts in world-writable directories.

Step 2: Detection — Query endpoint and SIEM telemetry for unexpected loading of NSS modules by kernel helper processes (look for '/sbin/request-key' or 'cifs.upcall' spawning unusual child processes or loading shared libraries from non-standard paths). Monitor audit logs for privilege escalation events: ausearch -m USER_AUTH,PRIV_ESC,SYSCALL -ts recent. Enable auditd rules tracking execve and mmap syscalls by the cifs.upcall helper. Check for unusual entries in /etc/nsswitch.conf or unexpected .so files in NSS library paths. Reference: NIST AU-2 (Event Logging), NIST SI-4 (System Monitoring), CIS 8.2 (Collect Audit Logs).

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis

Controls: NIST AU-2 (Event Logging), NIST AU-12 (Audit Record Generation), NIST SI-4 (System Monitoring), NIST AU-6 (Audit Record Review, Analysis, and Reporting), CIS 8.2 (Collect Audit Logs)

Compensating: Deploy the following auditd rules immediately on all affected hosts to instrument the CIFSwitch attack path — add to /etc/audit/rules.d/cifswitch.rules: '-a always,exit -F exe=/sbin/request-key -S execve -k cifswitch_exec' and '-a always,exit -F exe=/sbin/cifs.upcall -S mmap -S open -k cifswitch_mmap'. Then run: 'auditctl -R /etc/audit/rules.d/cifswitch.rules && augenrules --load'. Query existing logs with: 'ausearch -k cifswitch_exec -ts today' and 'ausearch -m SYSCALL -sc mmap -ts today | grep cifs'. For NSS injection detection without SIEM, use: 'inotifywait -m /lib/x86_64-linux-gnu/ /usr/lib64/ -e create,modify --include "libnss_.*\.so" &' to watch for new NSS library drops in real time.

Evidence: Capture before enabling new auditd rules to preserve pre-detection state: (1) 'ausearch -m SYSCALL,EXECVE -ts boot -te now > /tmp/prescan_syscall_audit.log' — the CIFSwitch exploit triggers execve from within the cifs.upcall/request-key context to execute a root shell, which will appear as an anomalous parent-child process relationship in SYSCALL records; (2) 'journalctl -u cifs.upcall --since boot' and 'journalctl -u request-key --since boot' to retrieve any prior invocations that could indicate exploitation attempts; (3) 'cat /proc/*/maps 2>/dev/null | grep -E "(cifs|nss).*rwx"' to identify any currently running process with a writable+executable NSS or CIFS-related memory mapping, which is a direct indicator of in-progress exploit execution; (4) 'ps auxf | grep -E "(cifs.upcall|request-key)"' to document current process tree state before any remediation.

Step 3: Eradication — Apply vendor-issued patched kernels as they become available. AlmaLinux has published patched kernels for community testing as of 2026-05-28 (see https://almalinux.org/blog/2026-05-28-cifswitch/); monitor CentOS Stream, Rocky Linux, SUSE, and upstream kernel channels for corresponding fixes. If cifs-utils is not required for operations, remove it permanently. If

CIFS mounts are required, restrict user namespace creation where feasible: set 'kernel.unprivileged_userns_clone=0' (Debian-family) or 'user.max_user_namespaces=0' (RHEL-family) via sysctl as a temporary mitigation — validate operational impact before deploying. Reference: NIST SI-2 (Flaw Remediation), CIS 7.3 (Perform Automated OS Patch Management).

NIST Phase: Eradication

Reference: NIST 800-61r3 §3.4 — Eradication

Controls: NIST SI-2 (Flaw Remediation), NIST CM-6 (Configuration Settings), NIST CM-7 (Least Functionality), CIS 7.3 (Perform Automated Operating System Patch Management), CIS 7.4 (Perform Automated Application Patch Management), CIS 4.6 (Securely Manage Enterprise Assets and Software)

Compensating: For teams without automated patch orchestration: use 'dnf update kernel cifs-utils --advisory=$(dnf advisory list | grep -i cifs)' on RHEL-family hosts once advisories are published, or track updates via 'dnf updateinfo list security' daily. Apply the user namespace sysctl mitigation immediately without waiting for kernel patches — it directly blocks the unprivileged namespace creation that the CIFSwitch exploit requires: 'echo "user.max_user_namespaces=0" >> /etc/sysctl.d/99-cifswitch-mitigation.conf && sysctl -p /etc/sysctl.d/99-cifswitch-mitigation.conf'. Verify the mitigation applied: 'sysctl user.max_user_namespaces' must return 0. Note: this will break rootless Podman/Docker and Flatpak — test on a non-production host first and document the operational exception if those services are in use.

Evidence: Before applying patches or sysctl changes, preserve: (1) 'uname -r && rpm -q kernel cifs-utils' (or 'dpkg -l linux-image-* cifs-utils') — document the exact vulnerable kernel and cifs-utils versions for regulatory evidence and change records; (2) 'sysctl -a | grep -E "(userns|user_namespaces|unprivileged)" > /tmp/pre_patch_sysctl_baseline.txt' to establish the pre-mitigation namespace configuration; (3) if any exploitation is suspected prior to eradication, capture a full memory image of affected hosts using LiME (Linux Memory Extractor) before rebooting into the patched kernel, as the post-reboot memory state will not contain exploit artifacts from the vulnerable kernel session; (4) 'cat /proc/sys/user/max_user_namespaces' and 'cat /proc/sys/kernel/unprivileged_userns_clone' to document the attack-surface-enabling configuration state that persisted during the vulnerable window.

Step 4: Recovery — After patching, verify the installed kernel version matches the vendor-issued fix and that cifs-utils has been updated or removed as appropriate. Re-enable any services paused during containment. Run a post-remediation audit: confirm no unauthorized accounts were created, no cron jobs or startup scripts were modified, and no new SUID binaries exist (find / -perm -4000 -type f 2>/dev/null). Review auditd logs for any exploitation activity that occurred before patching. Reference: NIST AU-6 (Audit Record Review), NIST IR-4 (Incident Handling), D3-SFA (System File Analysis), D3-SICA (System Init Config Analysis).

NIST Phase: Recovery

Reference: NIST 800-61r3 §3.5 — Recovery

Controls: NIST IR-4 (Incident Handling), NIST AU-6 (Audit Record Review, Analysis, and Reporting), NIST CM-6 (Configuration Settings), NIST SI-2 (Flaw Remediation), CIS 5.1 (Establish and Maintain an Inventory of Accounts), CIS 7.1 (Establish and Maintain a Vulnerability Management Process)

Compensating: Execute a structured post-patch verification script on each remediated host: (1) 'uname -r' — confirm kernel version matches vendor advisory; (2) 'rpm -q cifs-utils || dpkg -l cifs-utils' — confirm patched or removed; (3) 'awk -F: "$3 == 0" /etc/passwd' — list all UID 0 accounts (any entry other than root is a persistence indicator from successful CIFSwitch exploitation where attacker escalated to root and created a backdoor account); (4) 'find / -perm -4000 -type f -newer /var/log/audit/audit.log.1 2>/dev/null' — find SUID binaries created or modified since the last log rotation, which would indicate attacker persistence via a SUID shell backdoor installed after root escalation via CIFSwitch; (5) 'crontab -l -u root && ls -la /etc/cron* /var/spool/cron/' — detect root cron persistence.

Evidence: Before re-enabling paused services, collect for the incident record: (1) 'ausearch -k cifswitch_exec -k cifswitch_mmap -ts boot -te now' — retrieve all hits on the auditd rules deployed in Step 2 to determine if exploitation occurred in the pre-patch window; (2) 'find /etc/nsswitch.conf /lib/x86_64-linux-gnu/libnss_*.so* /usr/lib64/libnss_*.so* -newer /var/log/audit/audit.log.1' — identify any NSS configuration or library files modified during the vulnerable period, which is the specific persistence mechanism relevant to CIFSwitch's NSS-hijacking attack path; (3) 'last -F | head -50' and 'lastb -F | head -50' — reconstruct login history across the vulnerable window to correlate any root session with

exploit timing; (4) '/var/log/secure' (RHEL-family) or '/var/log/auth.log' (Debian-family) grep for 'su\|sudo\|COMMAND.*root' entries from non-admin accounts during the exposure window.

Step 5: Post-Incident — Conduct a control gap review: assess whether least-privilege policies prevented unnecessary local user access to affected hosts (NIST AC-6), whether software inventory practices would have flagged cifs-utils as an unnecessary installed package (CIS 2.1 — Establish and Maintain a Software Inventory), and whether user namespace restrictions are standardized in hardening baselines (CIS 4.6 — Securely Manage Enterprise Assets and Software). Update vulnerability management SLAs to account for unassigned-CVE threats with public PoC exploits. Reference: NIST CM-6 (Configuration Settings), CIS 7.1 (Establish and Maintain a Vulnerability Management Process).

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity

Controls: NIST AC-6 (Least Privilege), NIST CM-6 (Configuration Settings), NIST RA-3 (Risk Assessment), NIST SI-2 (Flaw Remediation), CIS 2.1 (Establish and Maintain a Software Inventory), CIS 4.6 (Securely Manage Enterprise Assets and Software), CIS 7.1 (Establish and Maintain a Vulnerability Management Process), CIS 7.2 (Establish and Maintain a Remediation Process)

Compensating: Operationalize two durable controls to prevent recurrence of this class of threat: (1) Add a cifs-utils presence check to your standard Linux build audit — encode 'rpm -q cifs-utils' (RHEL) or 'dpkg -l cifs-utils' (Debian) into your weekly osquery scheduled query ('SELECT name, version FROM deb_packages WHERE name = "cifs-utils"') so any re-installation triggers an alert; (2) Encode the user namespace restriction as a mandatory CIS benchmark check — add 'user.max_user_namespaces=0' to your sysctl hardening template (Ansible role, kickstart %post, or cloud-init) and validate compliance with: 'ansible all -m command -a "sysctl user.max_user_namespaces" | grep -v ": 0"' to flag any host that reverted or was deployed without the control. Document the CIFSwitch-specific rationale in the hardening baseline so future teams understand the threat context driving the control.

Evidence: Compile for the lessons-learned record and potential regulatory reporting: (1) Timeline of exposure window — from cifs-utils 6.14 installation date (from rpm/dpkg install timestamps) to patch application date — this defines the window during which unauthorized root access was possible on each host and is required for any breach notification assessment; (2) Full list of user accounts with local interactive access (shell != /sbin/nologin or /bin/false) on affected hosts during the exposure window — these are the accounts that could have exploited CIFSwitch, and their activity during that period requires review; (3) Consolidated auditd logs from all affected hosts covering the exposure window, specifically SYSCALL records for execve events originating from cifs.upcall or request-key process contexts, retained per NIST AU-11 (Audit Record Retention) requirements; (4) Pre- and post-remediation SUID binary inventory diff to establish whether any persistence mechanism was installed via root access gained through CIFSwitch exploitation.

Detection Guidance

Primary detection focus: unexpected execution of privileged kernel helper processes loading attacker-supplied shared libraries. Specific indicators to query: (1) Audit log events showing '/sbin/request-key' or '/usr/sbin/cifs.upcall' spawning child processes or executing binaries outside expected paths, use 'ausearch -m EXECVE' filtered by those parent process names. (2) NSS module loads from non-standard paths: monitor open/openat syscalls by cifs.upcall or request-key for .so files outside /lib, /usr/lib, or vendor-expected NSS directories. (3) Unexpected privilege transitions: audit records with PRIV_ESC or USER_AUTH events following CIFS-related kernel upcalls. (4) New SUID binaries or modifications to /etc/nsswitch.conf, baseline and alert on deviations. (5) Unusual processes running as root spawned from low-privilege user sessions. SIEM rule focus: correlate auditd SYSCALL records for mmap/execve by cifs-related helpers with subsequent UID=0 process creation. No public network IOCs (IP, domain, hash) have been reported for this vulnerability; exploitation is local and leaves OS-level artifacts, not network signatures. Reference: NIST AU-2, NIST AU-12, NIST SI-4, CIS 8.2, NIST AU-2 (Audit Events), NIST SI-3 (Malicious Code Protection).

Framework Mappings

MITRE-ATTACK

- **T1068** — Exploitation for Privilege Escalation
- **T1055** — Process Injection
- **T1574.006** — Dynamic Linker Hijacking
- **T1548.001** — Setuid and Setgid

NIST-800-53R5

- **AC-6** — Least Privilege
- **SC-7** — Boundary Protection
- **SI-2** — Flaw Remediation
- **SI-3** — Malicious Code Protection
- **SI-4** — System Monitoring
- **AC-3** — Access Enforcement

OWASP-TOP10-2021

- **A01:2021** — Broken Access Control

CIS-V8

- **5.4** — Restrict Administrator Privileges to Dedicated Administrator Accounts
- **6.8** — Define and Maintain Role-Based Access Control
- **6.1** — Establish an Access Granting Process
- **6.2** — Establish an Access Revoking Process
- **7.3** — Perform Automated Operating System Patch Management
- **7.4** — Perform Automated Application Patch Management

SOC2-TSC

- **CC6.1** — The entity implements logical access security software, infrastructure, and architectures over protected information assets
- **CC6.3** — Authorizes, modifies, or removes access

HIPAA-SECURITY

- **164.312(a)(1)** — Access Control

ISO-27001-2022

- **A.8.8** — Management of technical vulnerabilities

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1068	Exploitation for Privilege Escalation	Privilege-Escalation
T1055	Process Injection	Defense-Evasion
T1574.006	Dynamic Linker Hijacking	Persistence
T1548.001	Setuid and Setgid	Privilege-Escalation

Sources

Source	URL	Tier
Security News	https://www.bleepingcomputer.com/news/security/new-cifswitch-linux-...	T3
CIFSwitch: Help Us Test the Patched Kernels - AlmaLinux	https://almalinux.org/blog/2026-05-28-cifswitch/	T3
New Linux CIFSwitch Kernel Vulnerability Allows Attackers to Gain ...	https://www.reddit.com/r/linux/comments/1tqgk9a/new_linux_cifswitch...	T3
CIFSwitch flaw hands root to any Linux local user - AI Weekly	https://aiweekly.co/alerts/cifswitch-flaw-hands-root-to-any-linux-l...	T3
CIFSwitch Linux Kernel Flaw Grants Local Root on cifs-utils - TuxCare	https://tuxcare.com/de/blog/cifswitch-cve/	T3

Generated 2026-05-30 19:07 UTC by TJS Security Command Center