INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-06-04 19:19 UTC

Agentic AI Attack Surface Confirmed: Red Team Data Validates 7 New Failure Modes as Zero-Click Exploit Chains Emerge

SECURITY ANALYSIS | CRITICAL | CVSS 9.5

SCC Item ID	SCC-STY-2026-0167
Type	Security Analysis
CVE ID	CVE-2026-25253
Severity	CRITICAL
CVSS Base Score	9.5
Affected Products	Microsoft Security Copilot, OpenClaw agentic framework, Model Context Protocol (MCP) ecosystem, agentic AI systems broadly
Published	2026-06-04T19:14:42+00:00
Discovery Source	Rss:T1 Threatintel

Executive Summary

Microsoft's AI Red Team has formally documented seven new failure modes in agentic AI systems after 12 months of production red team engagements, confirming that zero-click exploit chains can achieve data exfiltration and lateral movement from external inputs alone, no user interaction required beyond initial agent deployment. Simultaneously, 31 commercially operating groups have been identified deploying AI Recommendation Poisoning at scale, targeting agent memory and context stores, while 336 malicious plugins were confirmed in the OpenClaw marketplace. This represents a shift in the enterprise threat surface: agentic AI systems are now confirmed exploitable infrastructure, and organizations deploying agentic systems without embedded safety evaluation or trust boundary controls are operating with an unvalidated attack surface.

Technical Analysis

Microsoft's AI Red Team v2.0 taxonomy, grounded in 12 months of engagements against production agentic systems including Microsoft Security Copilot and the Model Context Protocol (MCP) ecosystem, formalizes seven new failure modes: trust boundary violations, tool call injection, insecure memory handling, and four additional categories not fully enumerated in secondary source coverage. The most exploited failure mode confirmed across engagements is human-in-the-loop (HITL) bypass, the ability to trigger consequential agent

actions without a human approval step being invoked. The attack chain documented achieves end-to-end exploitation from external input: an adversary crafts a malicious prompt or poisoned data object that the agent ingests, the agent executes tool calls or API actions in response, and data exfiltration or lateral movement occurs entirely within the agent's authorized capability scope. No user beyond the initial deployer is involved.

In parallel, Microsoft's February 2026 research documents a distinct but related threat: adversaries injecting malicious content into AI memory stores and retrieval-augmented generation (RAG) pipelines to persistently influence agent recommendations and decisions. Commercial deployment of this technique has been identified across multiple organizations, illustrating the blurred boundary between manipulation-as-a-service and active exploitation.

The OpenClaw agentic framework marketplace presents a supply chain dimension: confirmed malicious plugins represent a plugin ecosystem compromise vector (MITRE T1195) through which adversaries can introduce tool call injection payloads or insecure memory handling behaviors into otherwise legitimate workflows. This maps directly to MITRE ATT&CK for ML technique AML.T0051 (LLM Prompt Injection) and AML.T0080 (Memory Poisoning), as well as T1195 (Supply Chain Compromise) and T1190 (Exploit Public-Facing Application) for MCP-exposed endpoints.

Multiple security organizations' research independently validates the need for continuous, embedded safety evaluation across multi-component workflows. The core defensive gap exploited across all documented scenarios is the absence of runtime trust enforcement between agent components, agents inheriting excessive permissions, tool calls executing without scoped authorization, and memory stores accepting unvalidated writes. CVE-2026-25253 has been assigned CVSS 9.5 (Critical) and appears to map to implementation flaws within this threat surface. However, the specific vulnerability details require validation against NVD and MSRC primary records, as secondary sources do not yet provide full enumeration. Verify patch and mitigation status directly before deployment decisions.

Action Checklist

1. Step 1: Assess exposure, inventory all agentic AI deployments in your environment, including Microsoft Security Copilot, any MCP-integrated systems, OpenClaw framework instances, and third-party agentic tools; identify which have external-facing inputs or operate with broad tool-call permissions
2. Step 2: Review controls, audit HITL enforcement configurations on all deployed agents; verify that agent-to-tool and agent-to-agent calls are scoped by least privilege (NIST AC-6) and that memory/context stores reject unvalidated external writes; apply CIS 3.3 (Configure Data Access Control Lists) to RAG and memory store backends
3. Step 3: Audit plugin and MCP ecosystem hygiene, cross-reference installed OpenClaw plugins and MCP connectors against known malicious plugins; remove unauthorized software per CIS 2.3 (Address Unauthorized Software) and enforce allowlisting via CIS 2.1 (Establish and Maintain a Software Inventory)
4. Step 4: Implement continuous safety evaluation, deploy runtime monitoring across multi-component agentic workflows per vendor safety frameworks; instrument tool call logs for anomalous execution patterns consistent with AML.T0051 (LLM Prompt Injection) and AML.T0080 (Memory Poisoning); align logging to NIST AU-2 (Event Logging) and AU-12 (Audit Record Generation)
5. Step 5: Update threat model, formally add trust boundary violation, tool call injection, insecure memory handling, and HITL bypass to your threat register; map to MITRE ATT&CK techniques T1195, T1190, T1059, AML.T0051, AML.T0080, and T1557; assign ownership for each failure mode

6. Step 6: Validate CVE-2026-25253 patch status, check NVD (https://nvd.nist.gov/vuln/detail/CVE-2026-25253) and MSRC (https://msrc.microsoft.com/update-guide/vulnerability/CVE-2026-25253) directly for current patch availability. Note: Secondary sources indicate this CVE details may not yet be fully indexed; contact Microsoft support directly if NVD record is incomplete
7. Step 7: Communicate findings, brief leadership on the structural risk shift: agentic AI systems are now exploitable infrastructure, not productivity tools with soft risk; present specific exposure based on your deployment inventory, not generic AI risk language

IR / Forensic Enrichment

Triage Priority	IMMEDIATE
Escalation Criteria	Escalate immediately to CISO and legal counsel if the Step 3 plugin audit identifies any of the 336 confirmed malicious MCP plugins installed in your environment, if Security Copilot or OpenClaw audit logs show tool-call sequences consistent with AML.T0051 prompt injection or AML.T0080 memory poisoning (external input → memory write → external API call chain within a single agent session), or if the organization is subject to HIPAA, PCI-DSS, or SEC cyber disclosure rules and any agent with access to regulated data cannot be confirmed as uncompromised before patch availability.
Recovery Notes	After patching CVE-2026-25253 and removing confirmed malicious MCP plugins, flush and rebuild all agent memory and RAG context stores from validated clean sources — do not assume memory store integrity after potential AML.T0080 exploitation, as poisoned context persists across agent restarts and patch cycles. Re-enable external-facing MCP connectors only after confirming patched versions, reimposing least-privilege tool-call scopes, and validating HITL enforcement is active for all high-risk tool categories (file write, external API call, credential access). Monitor Security Copilot Unified Audit Logs and OpenClaw tool_call_audit.log continuously for 30 days post-recovery, specifically watching for the re-emergence of external-origin memory writes or unexpected tool-call sequences that would indicate persistence mechanisms survived eradication.

Forensic Artifacts	Microsoft Purview Unified Audit Log — filter on 'CopilotInteraction' and 'AIPluginInvocation' operation types for the Security Copilot tenant; look for tool invocations where the input source traces to an external MCP connector and the output destination is an external API endpoint, which is the forensic signature of the zero-click exfiltration chain documented by Microsoft's AI Red Team | OpenClaw plugin execution log (/var/log/openclaw/plugin_exec.log or ~/.openclaw/logs/) — extract entries showing plugin name, invocation trigger (internal vs. external input), declared vs. actual tool calls made, and any outbound network connections; mismatches between declared plugin behavior and actual tool calls are the primary indicator of a malicious plugin from the 336 confirmed set | Agent memory and RAG store write audit trail — query the vector database or key-value store backend (Chroma, Qdrant, Redis) for write operations where the data origin is flagged as external or unvalidated; AML.T0080 Memory Poisoning leaves a specific artifact pattern of externally-sourced content written to agent context without sanitization, retrievable via the store's native audit or transaction log | MCP connector network traffic capture — PCAP or netflow records showing outbound connections from agent host processes to non-inventoried external endpoints during agent task execution windows; correlate connection timestamps against agent tool-call log timestamps to identify whether data left the environment during exploit chain execution | OpenClaw plugin binary hashes and manifest files — SHA-256 hashes of all installed plugin .so or .whl files compared against the malicious plugin IOC list; additionally inspect plugin manifest declared_permissions vs. runtime_permissions fields for privilege escalation indicators where a plugin claims limited scope but executed broader tool calls

Per-Action IR Details

Step 1: Assess exposure — inventory all agentic AI deployments in your environment, including Microsoft Security Copilot, any MCP-integrated systems, OpenClaw framework instances, and third-party agentic tools; identify which have external-facing inputs or operate with broad tool-call permissions

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Establishing IR capability requires knowing what agentic AI infrastructure exists, which external-facing inputs are active, and which agents operate with broad tool-call permissions that could serve as zero-click exploit entry points

Controls: CIS 1.1 (IG1/IG2/IG3) — Establish and Maintain Detailed Enterprise Asset Inventory, CIS 2.1 (IG1/IG2/IG3) — Establish and Maintain a Software Inventory, NIST AC-20 — Use Of External Systems

Compensating: Run 'Get-Command *Copilot* | Export-Csv copilot_inventory.csv' in PowerShell to surface Security Copilot plugin registrations; enumerate MCP connector configs via 'find / -name mcp_config.json 2>/dev/null' on Linux hosts running OpenClaw; document each agent's tool-call permission scope in a spreadsheet, flagging any agent with write access to external APIs or memory stores as HIGH EXPOSURE.

Evidence: Before inventorying, snapshot the current state: export Microsoft Security Copilot plugin registry via M365 Admin Center > Integrated Apps; collect OpenClaw framework config files (typically openclawrc.yaml or .openclawconfig) which enumerate connected MCP endpoints and declared tool permissions; capture current memory/context store backend connection strings from agent deployment manifests — these establish the pre-incident baseline and may reveal already-compromised configurations.

Step 2: Review controls — audit HITL enforcement configurations on all deployed agents; verify that agent-to-tool and agent-to-agent calls are scoped by least privilege (NIST AC-6) and that memory/context stores reject unvalidated external writes; apply CIS 3.3 (Configure Data Access Control Lists) to RAG and memory store backends

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Hardening agent trust boundaries and HITL enforcement before an active exploit chain is triggered is foundational preparation; zero-click chains documented in the Microsoft AI Red

Team findings require no user interaction, making pre-deployment control validation the only preventive window

Controls: NIST AC-6 — Least Privilege, NIST AC-3 — Access Enforcement, CIS 3.3 (IG1/IG2/IG3) — Configure Data Access Control Lists

Compensating: For RAG/vector store backends (e.g., Chroma, Weaviate, Qdrant running locally): review the authentication config file and confirm no unauthenticated write endpoints are exposed — run 'curl -X POST http://localhost:8080/v1/objects' and confirm a 401/403 response; for OpenClaw agents, inspect tool_permissions blocks in agent manifests and remove any wildcard (*) tool-call grants; document every HITL bypass exception with a written justification and owner.

Evidence: Capture HITL configuration state before modifying: export agent policy configs showing current human-approval thresholds (file paths vary by framework — for Security Copilot, pull from M365 Defender > Settings > Microsoft Copilot for Security > Plugin permissions); for OpenClaw, extract agent_policy.json or equivalent; query the RAG/memory store's access control list to identify any write permissions granted to external-origin data without validation rules — this establishes whether insecure memory handling (one of the seven documented failure modes) is already present.

Step 3: Audit plugin and MCP ecosystem hygiene — cross-reference installed OpenClaw plugins and MCP connectors against the 336 confirmed malicious plugins; remove unauthorized software per CIS 2.3 (Address Unauthorized Software) and enforce allowlisting via CIS 2.1 (Establish and Maintain a Software Inventory)

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection & Analysis: Cross-referencing installed MCP plugins against the 336 confirmed malicious plugin list is an active threat detection action — identifying a match constitutes confirmation of a potential supply chain compromise requiring immediate escalation to containment

Controls: CIS 2.1 (IG1/IG2/IG3) — Establish and Maintain a Software Inventory, CIS 2.2 (IG1/IG2/IG3) — Ensure Authorized Software is Currently Supported, CIS 2.3 (IG1/IG2/IG3) — Address Unauthorized Software

Compensating: Extract installed OpenClaw plugin manifests ('openclawctl plugin list --json > installed_plugins.json') and MCP connector registrations; compute SHA-256 hashes of each plugin binary ('sha256sum /opt/openclaw/plugins/*.so > plugin_hashes.txt'); compare plugin names, publisher IDs, and hashes against the 336 malicious plugin IOC list released alongside this advisory; flag any match as a Priority 1 removal and preserve the binary for forensic analysis before deletion.

Evidence: Before removing any plugin: capture the full plugin installation directory with timestamps ('ls -la /opt/openclaw/plugins/'); extract plugin manifest metadata including declared permissions, external call endpoints, and version strings; review OpenClaw plugin execution logs (typically ~/.openclaw/logs/plugin_exec.log or /var/log/openclaw/) for evidence of data exfiltration callbacks — specifically look for outbound POST requests to non-inventory endpoints made during agent task execution, which would indicate active AI Recommendation Poisoning exploitation by one of the 31 identified commercial threat groups.

Step 4: Implement continuous safety evaluation — deploy runtime monitoring across multi-component agentic workflows per NVIDIA and Lakera guidance; instrument tool call logs for anomalous execution patterns consistent with AML.T0051 (LLM Prompt Injection) and AML.T0080 (Memory Poisoning); align logging to NIST AU-2 (Event Logging) and AU-12 (Audit Record Generation)

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection & Analysis: Runtime instrumentation of agent tool-call chains enables real-time detection of the zero-click exploit patterns documented by Microsoft's AI Red Team, where external inputs traverse agent memory and tool-call boundaries without user interaction; NIST 800-61r3 DE.CM-09 specifically calls for monitoring common attack vectors including the agent-mediated channels that AML.T0051 and AML.T0080 exploit

Controls: NIST AU-2 — Event Logging, NIST AU-12 — Audit Record Generation, NIST AU-6 — Audit Record Review, Analysis, And Reporting

Compensating: Deploy Sysmon on Windows hosts running Security Copilot or OpenClaw agents with EventID 1 (Process Create) and EventID 3 (Network Connection) rules to capture unexpected child processes or outbound connections spawned during agent tool calls; write a Sigma rule targeting agent process trees where a network connection follows an unexpected tool invocation sequence (e.g., memory read → external POST within 500ms); for

Linux OpenClaw deployments, use auditd rules ('auditctl -a always,exit -F arch=b64 -S execve -k agent_exec') to log all process execution within the agent runtime user context.

Evidence: Before deploying new monitoring, extract the current baseline of agent tool-call logs to identify pre-existing anomalies: for Security Copilot, pull audit logs from Microsoft Purview Audit (search for 'CopilotInteraction' and 'AIPluginInvocation' operations in the Unified Audit Log); for OpenClaw, collect tool_call_audit.log entries showing tool name, input source (internal vs. external), output destination, and timestamp; capture any existing memory store write events that originated from external agent inputs — these are the primary forensic indicator of AML.T0080 Memory Poisoning activity predating your monitoring deployment.

Step 5: Update threat model — formally add trust boundary violation, tool call injection, insecure memory handling, and HITL bypass to your threat register; map to MITRE ATT&CK techniques T1195, T1190, T1059, AML.T0051, AML.T0080, and T1557; assign ownership for each failure mode

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: Updating the threat register with the seven formally documented agentic AI failure modes constitutes the lessons-learned and threat model improvement function; even without a confirmed incident, the Microsoft AI Red Team's 12-month production red team findings constitute validated intelligence requiring formal incorporation into the organization's risk register

Controls: NIST AC-4 — Information Flow Enforcement, NIST AC-5 — Separation Of Duties

Compensating: Document each of the seven failure modes in a structured threat register entry using a free template (OWASP Threat Dragon or a spreadsheet): columns for failure mode name, ATT&CK technique ID, affected product (Security Copilot / OpenClaw / MCP ecosystem), current control status (mitigated/partial/unmitigated), and assigned owner; conduct a 30-minute tabletop using the zero-click exploit chain scenario — external prompt injection via MCP connector → memory poisoning → lateral tool call → data exfiltration — to validate that each failure mode has a named owner and a detection hypothesis.

Evidence: Before finalizing the threat model update, gather intelligence inputs that should inform it: collect the Microsoft AI Red Team's formal failure mode documentation (reference the advisory associated with CVE-2026-25253); extract your current agent deployment inventory from Step 1 to map each failure mode against specific deployed systems; pull existing SIEM or audit log data to determine whether any of the seven failure modes have already generated observable signals in your environment — anomalous tool call sequences, unexpected memory store writes, or HITL approval rate anomalies should be documented as pre-existing indicators before they are classified as new threats.

Step 6: Validate CVE-2026-25253 patch status — check NVD (https://nvd.nist.gov/vuln/detail/CVE-2026-25253) and MSRC (https://msrc.microsoft.com/update-guide/vulnerability/CVE-2026-25253) directly for current patch availability and compensating controls; do not rely on secondary RSS sources for remediation guidance on a CVSS 9.5 finding

NIST Phase: Containment

Reference: NIST 800-61r3 §3.3 — Containment Strategy: For a CVSS 9.5 zero-click exploit chain affecting Microsoft Security Copilot and the MCP ecosystem, patch validation against authoritative vendor sources (MSRC and NVD) is the immediate containment decision gate — acting on incomplete or secondary-source patch information for a finding of this severity risks both failed containment and operational disruption

Controls: CIS 7.1 (IG1/IG2/IG3) — Establish and Maintain a Vulnerability Management Process, CIS 7.2 (IG1/IG2/IG3) — Establish and Maintain a Remediation Process, CIS 7.3 (IG1/IG2/IG3) — Perform Automated Operating System Patch Management, CIS 7.4 (IG1/IG2/IG3) — Perform Automated Application Patch Management

Compensating: For Security Copilot: check patch status via Microsoft 365 Admin Center > Health > Message Center filtering on CVE-2026-25253 or the associated KB article from MSRC; for OpenClaw: run 'openclawctl version --check-updates' and compare against the patched version listed in the vendor advisory; if no patch is available yet, implement the compensating control of disabling external-facing MCP connectors entirely ('openclawctl connector disable --all-external') until a patch is confirmed — this breaks the zero-click exploit chain's external input vector.

Evidence: Before applying any patch, snapshot the pre-patch state of affected systems: record current Security Copilot version from M365 Admin Center > Settings > Integrated Apps; capture OpenClaw framework version

('openclawctl version > pre_patch_version.txt'); export the current MCP connector configuration and plugin list with timestamps; take a memory dump or process snapshot of any actively running OpenClaw agent processes if forensic analysis of potential prior exploitation is required — patches may overwrite artifacts needed to determine whether exploitation occurred before remediation.

Step 7: Communicate findings — brief leadership on the structural risk shift: agentic AI systems are now exploitable infrastructure, not productivity tools with soft risk; present specific exposure based on your deployment inventory, not generic AI risk language

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: Leadership communication of the structural risk reclassification of agentic AI from productivity tooling to exploitable infrastructure is a formal post-incident improvement function; NIST 800-61r3 RS.MA-01 explicitly requires coordination with relevant stakeholders as part of incident response plan execution, including communicating scope and impact assessments to authorized personnel

Controls: NIST AC-1 — Policy And Procedures

Compensating: Build the leadership brief from your Step 1 inventory output: translate each agent deployment into a business-impact statement (e.g., 'Security Copilot has access to our M365 tenant data and can invoke external APIs — a zero-click exploit chain could exfiltrate tenant email and documents without analyst interaction'); include a one-page risk matrix mapping each of the seven documented failure modes against your specific deployed agents, their data access scope, and current control status; attach the MSRC advisory URL and CVSS 9.5 score as authority references rather than relying on internal framing alone.

Evidence: The communication package itself should be treated as a documented artifact: retain the leadership brief, the inventory data it was based on, and the date it was delivered as part of the incident record; if leadership decisions result in accepting residual risk (e.g., deferring patch deployment or keeping external MCP connectors active), capture that decision in writing with the approving authority's name — for a CVSS 9.5 finding with zero-click exploit chain confirmation, undocumented risk acceptance creates regulatory and liability exposure that must be preserved in the incident timeline.

Detection Guidance

Detection for agentic AI exploitation requires instrumentation at the agent runtime layer, not just the network perimeter. Priority log sources include agent orchestration logs (tool call sequences, memory read/write operations, external input ingestion events), MCP connector activity logs, and RAG pipeline query and retrieval logs.

Behavioral patterns to hunt for: (1) Tool call chains that execute in sequence without any logged human approval event, a signature of HITL bypass. (2) Memory store write operations originating from external-facing input channels (email ingestion, web browsing agents, document processors) rather than internal trusted sources, consistent with AML.T0080 (Memory Poisoning). (3) Outbound data transfers initiated by agent processes to destinations not in a pre-approved allowlist, potential exfiltration via T1530 or T1602. (4) Plugin or MCP connector registrations that post-date your last verified inventory, correlate against CIS 1.1 asset inventory and CIS 2.1 software inventory baselines. (5) Agents invoking scripting interpreters (T1059) or lateral movement tools (T1021) without a corresponding user-initiated workflow ticket.

For AI Recommendation Poisoning, audit agent-generated recommendations against a ground-truth baseline for unexpected vendor, product, or action recommendations, particularly in financial, procurement, or security workflow agents. Unexplained recommendation drift is a soft indicator.

Align audit log collection to NIST AU-2 and AU-12; ensure agent runtime events are retained per AU-11 for post-incident forensic reconstruction. D3FEND countermeasures with direct applicability: D3-UAP (User Account Permissions) for scoping agent tool-call authorizations; D3-LAM (Local Account Monitoring) for

detecting agent accounts performing anomalous actions; D3-SFA (System File Analysis) for monitoring memory store and context file integrity.

Indicators of Compromise

Type	Value	Context	Confidence
TOOL	`Pending — refer to Microsoft Security Blog (2026-06-04) for published indicators`	Tool call injection payloads and memory poisoning artifacts associated with zero-click agentic exploit chains; Microsoft AI Red Team report is the authoritative source for specific IOC values not reproduced in secondary coverage	LOW
TOOL	`Pending — refer to Microsoft Security Blog (2026-02-10) AI Recommendation Poisoning post for published indicators`	Memory store and RAG pipeline poisoning artifacts from 31 identified commercial operators deploying AI Recommendation Poisoning; specific injection payloads and behavioral signatures published in the source report	LOW
TOOL	`Pending — refer to OpenClaw marketplace security advisory for malicious plugin identifiers`	336 confirmed malicious plugins in the OpenClaw marketplace; plugin IDs, hashes, and behavioral signatures should be obtained directly from the OpenClaw security advisory or marketplace removal notices	LOW

Framework Mappings

MITRE-ATTACK

- **T1557** — Adversary-in-the-Middle
- **T1036** — Masquerading
- **T1195** — Supply Chain Compromise
- **T1602** — Data from Configuration Repository
- **T1190** — Exploit Public-Facing Application
- **T1059** — Command and Scripting Interpreter
- **T1021** — Remote Services
- **T1539** — Steal Web Session Cookie
- **T1078** — Valid Accounts
- **T1530** — Data from Cloud Storage
- **T1056** — Input Capture

NIST-800-53R5

- **SA-9** — External System Services

- **SR-2** — Supply Chain Risk Management Plan
- **SR-3** — Supply Chain Controls and Processes
- **SI-7** — Software, Firmware, and Information Integrity
- **CA-8** — Penetration Testing
- **RA-5** — Vulnerability Monitoring and Scanning
- **SC-7** — Boundary Protection
- **SI-2** — Flaw Remediation
- **CM-7** — Least Functionality
- **SI-3** — Malicious Code Protection
- **SI-4** — System Monitoring
- **AC-17** — Remote Access
- **AC-3** — Access Enforcement
- **IA-2** — Identification and Authentication (Organizational Users)
- **AC-2** — Account Management
- **AC-6** — Least Privilege
- **IA-5** — Authenticator Management
- **SI-10** — Information Input Validation
- **IA-8** — Identification and Authentication (Non-Organizational Users)

OWASP-TOP10-2021

- **A03:2021** — Injection
- **A07:2021** — Identification and Authentication Failures

CIS-V8

- **16.10** — Apply Secure Design Principles in Application Architectures
- **6.3** — Require MFA for Externally-Exposed Applications
- **6.4** — Require MFA for Remote Network Access
- **6.5** — Require MFA for Administrative Access

SOC2-TSC

- **CC6.1** — The entity implements logical access security software, infrastructure, and architectures over protected information assets
- **CC7.4** — Responds to identified security incidents

HIPAA-SECURITY

- **164.312(d)** — Person or Entity Authentication
- **164.308(a)(6)(ii)** — Response and Reporting

ISO-27001-2022

- **A.8.8** — Management of technical vulnerabilities
- **A.5.34** — Privacy and protection of personal information

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1557	Adversary-in-the-Middle	Credential-Access
T1036	Masquerading	Defense-Evasion
T1195	Supply Chain Compromise	Initial-Access
T1602	Data from Configuration Repository	Collection
T1190	Exploit Public-Facing Application	Initial-Access
T1059	Command and Scripting Interpreter	Execution
T1021	Remote Services	Lateral-Movement
T1539	Steal Web Session Cookie	Credential-Access
T1078	Valid Accounts	Defense-Evasion
T1530	Data from Cloud Storage	Collection
T1056	Input Capture	Collection

Sources

Source	URL	Tier
Microsoft Security Blog	https://www.microsoft.com/en-us/security/blog/2026/06/04/updating-t...	T1
	https://www.microsoft.com/en-us/security/blog/2026/06/04/updating-t...	T1
	https://www.helpnetsecurity.com/2025/12/08/nvidia-agentic-ai-securi...	T3
	https://www.microsoft.com/en-us/security/blog/2026/02/10/ai-recomme...	T1
CVE-2026-25253 detail - NVD	https://nvd.nist.gov/vuln/detail/CVE-2026-25253	T1
Microsoft Security Advisory	https://msrc.microsoft.com/update-guide/vulnerability/CVE-2026-25253	T1

Generated 2026-06-04 19:19 UTC by TJS Security Command Center