

INTELLIGENCE BRIEFING
Security Command Center

TLP:CLEAR
2026-05-14 13:50 UTC

Agentic AI Security Reaches Consensus: Five Eyes Agencies and Industry Align on Application-Layer Controls

GOVERNANCE | HIGH | CVSS 7.5

SCC Item ID	SCC-GOV-2026-0035
Type	Governance
Severity	HIGH
CVSS Base Score	7.5
Affected Products	Agentic AI frameworks broadly; Microsoft Security Copilot, Azure AI cited as reference implementations; no single vulnerable product identified
Published	2026-05-14T16:00:00+00:00
Discovery Source	Rss:T1 Threatintel

Executive Summary

Five Eyes intelligence agencies, CISA, NSA, Microsoft, and Palo Alto Networks have jointly established that autonomous AI agents introduce a new class of security risk requiring architectural controls at design time. Unlike traditional software vulnerabilities, agentic AI systems can be manipulated through crafted prompts to execute shell commands, escalate privileges, and operate beyond their intended scope, all using legitimate credentials and tools. Organizations deploying agentic AI without least-privilege identity controls, constrained tool permissions, and human approval gates for high-impact actions are accepting risk that existing software security frameworks were not designed to address.

Technical Analysis

Agentic AI frameworks inherit all existing software vulnerability classes (CWE-284, CWE-269, CWE-287, CWE-732, CWE-693) while introducing novel attack surfaces not addressed by conventional application security controls. Microsoft's May 7 2026 research documented concrete remote code execution paths (T1203, T1059) where prompt injection causes agent tool-invocation pipelines to execute arbitrary shell commands, prompts functioning as shell vectors. Additional confirmed attack paths include privilege escalation through agent identity gaps (T1548, T1134), agent identity spoofing and impersonation via weak or absent identity verification (T1078, CWE-287), unconstrained tool-use permissions enabling data manipulation (T1565) and defense impairment through audit log deletion (T1562), and exploitation of exposed agentic APIs (T1190). This guidance addresses architectural risk classes across agentic AI frameworks broadly; no single CVE is assigned. Microsoft Security

Copilot and Azure AI are cited as reference implementations with documented mitigations. The CISA/NSA joint publication, co-signed by AU, CA, NZ, and UK partner agencies, establishes agent identity, permissions, and escalation paths as first-class security primitives. Estimated impact severity (High) reflects high potential impact offset by exploitation complexity. No CISA KEV entry exists; this is a design-time governance item, not a discrete patchable vulnerability.

Action Checklist

1. Step 1: Inventory, catalog every agentic AI deployment in your environment: identify frameworks in use (LangChain, AutoGen, Azure AI Agent Service, Microsoft Security Copilot, and equivalents), their tool integrations, and the permissions each agent identity holds in production systems.
2. Step 2: Detection, audit agent identity logs for anomalous tool invocations, unexpected shell executions, permission scope changes, and outbound API calls inconsistent with defined agent purpose; review orchestration logs for prompt injection indicators such as instruction override patterns or unexpected task chains.
3. Step 3: Eradication, apply least-privilege agent identities: scope each agent's credentials and tool permissions to the minimum required for its defined function; remove standing permissions for high-impact actions (file write, code execution, external API calls) and replace with just-in-time grants requiring explicit authorization.
4. Step 4: Recovery, implement deterministic escalation gates for high-impact agent actions: require human-in-the-loop checkpoints before any agent executes shell commands, modifies access controls, deletes audit records, or makes external data transfers; validate that audit trails for agent actions are immutable and cannot be altered by the agent itself.
5. Step 5: Post-Incident, treat agent identity and permissions as part of your secure-by-design review process: update SDLC gates to require agentic AI security review against the CISA/NSA joint guidance and Microsoft's defense-in-depth framework before any agentic system reaches production; map agent capabilities to MITRE ATT&CK techniques relevant to your threat model.

IR / Forensic Enrichment

Triage Priority	URGENT
Escalation Criteria	Escalate immediately to CISO and legal if agent audit logs show tool invocations (shell execution, file write, external API calls) that cannot be attributed to a known authorized task — this indicates potential active exploitation of a prompt injection condition, which may constitute unauthorized access to production systems and could trigger breach notification obligations if the agent held access to PII, PHI, or regulated data.
Recovery Notes	After implementing least-privilege agent identities and human-in-the-loop gates, run each agent through its full intended task workflow under observation for a minimum of 5 business days, reviewing orchestration logs daily for any tool invocations that fall outside the newly defined permission scope — residual over-permissioned grants are common after partial remediation. Monitor Azure AD sign-in logs for the agent service principals for 30 days post-remediation to detect any credential reuse if API keys were not rotated during eradication. Verify immutable log configuration is enforced on each agent's log storage destination before returning the agent to unobserved production operation.

Forensic Artifacts	Azure AI Agent Service activity logs (Azure Monitor / Log Analytics workspace): filter on resource type 'Microsoft.MachineLearningServices' and 'Microsoft.CognitiveServices', specifically tool_call events with tool_name, tool_input, and tool_output fields — these are the primary record of what actions a manipulated agent actually executed LangChain verbose trace output or AutoGen conversation log files (default paths: ./logs/, ./autogen_logs/, or LANGCHAIN_TRACING_V2 destination): contain raw prompt, model response, and tool invocation chains that would reveal injected instructions and any instruction override patterns introduced via crafted user input or retrieved documents Azure AD sign-in logs filtered on agent service principal application IDs: timestamp, resource accessed, IP address, and conditional access result — anomalous resource access (outside the agent's declared target systems) or sign-ins from unexpected IPs indicate credential theft or lateral movement using the agent identity Microsoft Security Copilot audit logs (Microsoft 365 Compliance Center → Audit → SecurityCopilot workload): plugin invocation records, data connector queries, and session identifiers for any Copilot actions initiated outside of expected analyst working hours or by unexpected user principals — prompt injection in Copilot would appear as anomalous plugin calls in this log Outbound network flow logs from agent hosts (Azure NSG flow logs, VPC flow logs, or host-level tcpdump captures): destination IP/domain, byte count, and timing correlated with tool invocation timestamps — data exfiltration via a compromised agent would appear as high-volume or anomalous-destination outbound transfers tightly correlated with specific tool_call events in the orchestration logs
---------------------------	---

Per-Action IR Details

Step 1: Inventory — catalog every agentic AI deployment in your environment: identify frameworks in use (LangChain, AutoGen, Azure AI Agent Service, Microsoft Security Copilot, and equivalents), their tool integrations, and the permissions each agent identity holds in production systems.

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: establishing visibility into assets and identities before incidents occur; maps to CSF [GV, ID, PR] functions

Controls: NIST IR-4 (Incident Handling) — preparation pillar requires documented scope of systems under incident handling capability, NIST SI-5 (Security Alerts, Advisories, and Directives) — act on CISA/NSA joint guidance by identifying affected agentic deployments, NIST RA-2 (Security Categorization) — categorize each agent identity by the sensitivity of systems it can reach and actions it can invoke, CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory) — extend asset inventory to include agentic AI deployments, their framework versions, and service principal or API key identities, CIS 2.1 (Establish and Maintain a Software Inventory) — enumerate LangChain, AutoGen, Azure AI Agent Service SDK versions and associated tool-call plugins in the software inventory, CIS 5.1 (Establish and Maintain an Inventory of Accounts) — document every agent service principal, managed identity, OAuth client, or API key as a non-human account entry with its granted scopes

Compensating: Run `az ad sp list --all --query "[].{displayName:displayName, appId:appId}"` (Azure CLI) to enumerate service principals associated with AI workloads; cross-reference with LangChain/AutoGen config files (`.env`, `agent_config.yaml`, `tools.py`) checked into source repos using `grep -r 'tool' --include='*.py' --include='*.yaml'` to map declared tool permissions. For on-prem or multi-cloud, use osquery with `SELECT * FROM users; SELECT * FROM processes WHERE name LIKE '%agent%';` to surface running agent processes and their execution context. Maintain findings in a shared spreadsheet until a CMDB is available.

Evidence: Before inventorying live systems, snapshot current state to establish a pre-remediation baseline: export Azure AD audit logs (sign-in logs filtered on service principal logins) covering the past 90 days; capture the current permission grants for each agent managed identity via `az role assignment list --assignee --all`; preserve LangChain/AutoGen tool registration files and any `.env` files (redact secrets) documenting which shell tools, APIs, or file paths are exposed to the agent at runtime — these establish the blast radius if prompt injection has already occurred.

Step 2: Detection — audit agent identity logs for anomalous tool invocations, unexpected shell executions, permission scope changes, and outbound API calls inconsistent with defined agent purpose; review orchestration logs for prompt injection indicators such as instruction override patterns or unexpected task chains.

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis: correlating indicators of adverse events across multiple log sources to confirm incident scope; maps to CSF [DE] function

Controls: NIST SI-4 (System Monitoring) — monitor agentic AI orchestration layers for tool invocations and prompt content anomalies as a new monitoring surface, NIST AU-6 (Audit Record Review, Analysis, and Reporting) — review agent identity audit logs at defined frequency for instruction override patterns and unexpected task chains, NIST AU-2 (Event Logging) — confirm that agentic framework logging captures tool call name, arguments, response, and the originating prompt context for each invocation, NIST AU-3 (Content of Audit Records) — verify agent logs record who (agent identity), what (tool name and parameters), when (timestamp), where (target system or API endpoint), and why (task context), CIS 8.2 (Collect Audit Logs) — enable and centralize audit logging across Azure AI Agent Service activity logs, Microsoft Security Copilot audit events, and LangChain/AutoGen trace output

Compensating: For Azure-hosted agents, query Azure Monitor / Log Analytics: ``AzureActivity | where OperationNameValue contains 'Microsoft.MachineLearningServices' or OperationNameValue contains 'CognitiveServices' | where ActivityStatusValue == 'Success' | project TimeGenerated, Caller, OperationNameValue, Properties``. For LangChain/AutoGen running on Linux hosts, enable Sysmon (via `sysmonforlinux`) and filter on Event ID 1 (Process Create) where `ParentImage` matches the Python interpreter running the agent — unexpected ``bash``, ``sh``, ``curl``, or ``wget`` child processes are high-fidelity prompt injection indicators. Use the community Sigma rule ``proc_creation_win_susp_shell_spawn_from_interpreters.yml`` adapted for Python agent parent processes. For outbound API anomalies, capture agent host traffic with ``tcpdump -i eth0 -w agent_traffic.pcap 'port 443'`` and inspect with Wireshark, filtering on TLS SNI fields for domains outside the agent's declared API allow-list.

Evidence: Preserve before any containment action: full Azure AD sign-in logs for each agent service principal (portal: Azure AD → Sign-in logs → filter by application ID, export to JSON); Azure AI Agent Service activity logs showing `tool_call` events with their `tool_name`` and `tool_input`` fields; LangChain verbose trace output or AutoGen conversation logs stored in `./logs/`` or configured `autogen_logs/`` directory — these contain the raw prompt and model response chains that would reveal injected instructions; Microsoft Security Copilot audit logs (Microsoft 365 Compliance Center → Audit → filter on SecurityCopilot workload) for any plugin invocations or data connector queries initiated outside of expected analyst sessions; any outbound HTTP/S request logs from the agent host (proxy logs, NSG flow logs in Azure) showing destinations, volumes, and timing relative to anomalous tool calls.

Step 3: Eradication — apply least-privilege agent identities: scope each agent's credentials and tool permissions to the minimum required for its defined function; remove standing permissions for high-impact actions (file write, code execution, external API calls) and replace with just-in-time grants requiring explicit authorization.

NIST Phase: Eradication

Reference: NIST 800-61r3 §3.4 — Eradication: eliminating the conditions that enabled the incident, including over-permissioned identities that agentic AI attack chains exploit; maps to CSF [RS] function

Controls: NIST IR-4 (Incident Handling) — eradication actions must address the root enabling condition: standing high-impact permissions on agent identities that allow prompt injection to cause real damage, NIST AC-6 (Least Privilege) — restrict agent service principals and managed identities to only the permissions explicitly required for their defined automated function, NIST AC-2 (Account Management) — manage agent identities (service principals, managed identities, API keys) with the same lifecycle controls as human accounts, including periodic access review, NIST SI-2 (Flaw Remediation) — treat over-permissioned agent identity as a design flaw requiring remediation; apply vendor guidance from CISA/NSA joint advisory and Microsoft Azure AI security baseline, CIS 5.4 (Restrict Administrator Privileges to Dedicated Administrator Accounts) — agent identities must never hold persistent administrative roles; any elevated action must route through a separate privileged identity with explicit approval, CIS 6.1 (Establish an Access Granting Process) — apply formal access granting process to tool permission grants for agentic AI, requiring documented justification for each tool capability enabled

Compensating: In Azure, remove standing role assignments from agent managed identities: ``az role assignment delete --assignee --role Contributor --scope /subscriptions/`` and replace with narrowly scoped custom roles using ``az role definition create``. For LangChain/AutoGen, audit ``tools=[]`` lists in agent instantiation code and comment out or remove ``ShellTool``, ``PythonREPLTool``, ``FileWriteTool``, and any HTTP request tools not strictly required — commit the change through a pull request with mandatory review. For API key rotation, use ``az ad app credential reset --id`` to invalidate existing credentials and issue new scoped tokens. Document each permission removed and the business justification required to restore it, stored in the incident ticket as an audit trail per NIST AU-10 (Non-Repudiation).

Evidence: Before removing permissions, capture the pre-eradication permission state as forensic evidence of the over-permissioned condition: ``az role assignment list --assignee --all -o json > agent_permissions_pre_remediation.json``; export Azure AD app registration API permissions (portal → App registrations → agent app → API permissions → export); for LangChain/AutoGen, ``git log --follow -p agent_config.py`` to establish a history of tool additions as evidence of permission creep over time; preserve any Azure Policy compliance snapshots showing the agent identity was out of compliance with least-privilege policy — these records support post-incident reporting requirements under NIST IR-6 (Incident Reporting).

Step 4: Recovery — implement deterministic escalation gates for high-impact agent actions: require human-in-the-loop checkpoints before any agent executes shell commands, modifies access controls, deletes audit records, or makes external data transfers; validate that audit trails for agent actions are immutable and cannot be altered by the agent itself.

NIST Phase: Recovery

Reference: NIST 800-61r3 §3.5 — Recovery: restoring systems to a verified secure operational state with controls in place that prevent recurrence; maps to CSF [RC] function

Controls: NIST IR-4 (Incident Handling) — recovery must restore operations under conditions where the attack vector (unchecked autonomous agent action) is structurally blocked, not just patched, NIST AU-9 (Protection of Audit Information) — agent audit logs must be written to a destination the agent identity cannot access, modify, or delete; agent managed identity must have no write permissions on the log storage account or workspace, NIST AU-4 (Audit Storage Capacity) — allocate dedicated, agent-write-protected log storage to ensure agent action records are retained for post-incident and forensic use without risk of agent-driven deletion, NIST SI-7 (Software, Firmware, and Information Integrity) — implement integrity verification on agent orchestration configurations and tool allow-lists to detect unauthorized modification between deployments, NIST AC-5 (Separation of Duties) — the agent identity must not hold permissions to approve its own high-impact actions; human-in-the-loop gates must be enforced by a separate authorization service the agent cannot invoke or bypass, CIS 3.3 (Configure Data Access Control Lists) — apply explicit deny ACLs preventing agent managed identities from writing to or deleting from audit log storage containers

Compensating: For Azure AI Agent Service, implement human approval using Azure Logic Apps or Power Automate: trigger a Teams/email approval workflow when the agent requests a high-impact tool call, with the agent blocked on a polling loop until approval is received — this is achievable with free Azure Logic Apps consumption tier. For LangChain/AutoGen, wrap high-impact tools with a Python approval decorator that pauses execution and requires a CLI ``[y/N]`` confirmation from an operator before proceeding; commit this wrapper as a mandatory import in all agent instantiation modules. For immutable audit logs, configure Azure Storage immutability policies (``az storage container immutability-policy create``) on the log container, or write LangChain traces to a WORM-configured S3 bucket / append-only Elasticsearch index that the agent service account has no delete permissions on. Verify log immutability with ``az storage container immutability-policy show``.

Evidence: Before declaring recovery complete, collect: a post-remediation permission export (``az role assignment list --assignee --all -o json > agent_permissions_post_remediation.json``) to confirm least-privilege state; a test execution trace showing the human-in-the-loop gate firing correctly for a simulated high-impact tool call (preserve the trace log as evidence the gate works); Azure Storage immutability policy configuration export for the agent log container as evidence that audit records cannot be agent-altered; comparison diff between pre- and post-remediation tool allow-lists committed to source control, reviewed and approved by a second analyst — this diff constitutes the eradication evidence record referenced in NIST 800-61r3 §3.5.

Step 5: Post-Incident — treat agent identity and permissions as part of your secure-by-design review process: update SDLC gates to require agentic AI security review against the CISA/NSA joint guidance and Microsoft's

defense-in-depth framework before any agentic system reaches production; map agent capabilities to MITRE ATT&CK techniques relevant to your threat model.

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: lessons learned, policy updates, and detection improvements to prevent recurrence; maps to CSF [GV, ID] functions

Controls: NIST IR-4 (Incident Handling) — post-incident review must produce documented updates to the incident handling capability, specifically adding agentic AI as a named threat category in the IR plan, NIST IR-8 (Incident Response Plan) — update the IR plan to include agentic AI-specific playbooks covering prompt injection triage, agent identity isolation, and orchestration log collection procedures, NIST SI-2 (Flaw Remediation) — embed the CISA/NSA joint advisory on agentic AI security as a required security review checklist item in the SDLC gate for any system invoking LangChain, AutoGen, Azure AI Agent Service, or equivalent frameworks, NIST SA-11 (Developer Testing and Evaluation) — require security testing of agentic AI systems including adversarial prompt injection testing before production deployment, NIST RA-3 (Risk Assessment) — update organizational risk assessments to include agentic AI attack surface: prompt injection vectors, tool scope risks, and autonomous action blast radius, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — extend vulnerability management process to cover agentic AI framework updates (LangChain, AutoGen releases) and CISA/NSA advisory revisions as tracked vulnerability sources, CIS 7.2 (Establish and Maintain a Remediation Process) — define SLAs for remediating agentic AI misconfigurations (over-permissioned identities, missing human-in-the-loop gates) discovered through security reviews

Compensating: Create a one-page SDLC security review checklist for agentic AI derived directly from the CISA/NSA joint guidance (published at cisa.gov — verify current URL at time of use), covering: tool permission scope review, agent identity least-privilege verification, prompt injection input validation confirmation, human-in-the-loop gate implementation, and immutable audit log configuration. Store it in the team wiki and attach it as a required artifact to production deployment tickets. For ATT&CK mapping, use the free MITRE ATT&CK Navigator (attack.mitre.org/resources/attack-navigator/) to build a layer covering: T1059 (Command and Scripting Interpreter) for shell tool abuse, T1078 (Valid Accounts) for agent credential misuse, T1190 (Exploit Public-Facing Application) for prompt injection entry, T1530 (Data from Cloud Storage) for unauthorized exfiltration via agent API calls, and T1548 (Abuse Elevation Control Mechanism) for agent privilege escalation — export as JSON and store with the IR documentation.

Evidence: Document for the lessons-learned record: the complete agent inventory produced in Step 1 (baseline for future audits); the pre- and post-remediation permission diffs from Steps 3 and 4 (evidence of remediation completeness); detection query outputs from Step 2 showing the scope of anomalous tool invocations observed (confirms whether active exploitation occurred or risk was prospective); the ATT&CK Navigator layer export mapping this environment's agent capabilities to specific techniques (input to next threat hunt cycle); and a written lessons-learned memo documenting IR timeline, detection gaps, and SDLC process changes — this memo satisfies NIST IR-4 (Incident Handling) documentation requirements and feeds the next IR-3 (Incident Response Testing) exercise.

Detection Guidance

Monitor orchestration and tool-invocation logs for prompt injection indicators: look for instruction-override patterns ('ignore previous instructions', 'act as', 'you are now'), unexpected chaining of tool calls not consistent with the agent's defined workflow, and tool invocations that include shell metacharacters or command substitution syntax. In Azure AI and similar platforms, alert on agent identity tokens used outside expected resource scopes or time windows. Log and alert on any agent-initiated modifications to IAM policies, audit log configurations, or file ACLs. For Microsoft Security Copilot deployments, review plugin invocation logs for calls that were not user-initiated. Baseline normal agent behavior per deployment, volume of tool calls, typical target resources, expected execution duration, and alert on deviations. This guidance addresses architectural risk classes; detection is behavioral, not signature-based.

Framework Mappings

MITRE-ATTACK

- **T1548** — Abuse Elevation Control Mechanism
- **T1203** — Exploitation for Client Execution
- **T1059** — Command and Scripting Interpreter
- **T1190** — Exploit Public-Facing Application
- **T1078** — Valid Accounts
- **T1562** — Impair Defenses
- **T1134** — Access Token Manipulation
- **T1565** — Data Manipulation

NIST-800-53R5

- **AC-6** — Least Privilege
- **CM-6** — Configuration Settings
- **SC-7** — Boundary Protection
- **SI-2** — Flaw Remediation
- **SI-3** — Malicious Code Protection
- **SI-4** — System Monitoring
- **CM-7** — Least Functionality
- **SI-7** — Software, Firmware, and Information Integrity
- **CA-8** — Penetration Testing
- **RA-5** — Vulnerability Monitoring and Scanning
- **AC-2** — Account Management
- **IA-2** — Identification and Authentication (Organizational Users)
- **IA-5** — Authenticator Management
- **AU-9** — Protection of Audit Information
- **IA-8** — Identification and Authentication (Non-Organizational Users)
- **AC-3** — Access Enforcement
- **AT-2** — Literacy Training and Awareness

OWASP-TOP10-2021

- **A07:2021** — Identification and Authentication Failures
- **A01:2021** — Broken Access Control

CIS-V8

- **6.3** — Require MFA for Externally-Exposed Applications
- **6.4** — Require MFA for Remote Network Access
- **6.5** — Require MFA for Administrative Access
- **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
- **3.3** — Configure Data Access Control Lists
- **14.2** — Train Workforce Members to Recognize Social Engineering Attacks

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(d)** — Person or Entity Authentication
- **164.312(a)(1)** — Access Control
- **164.312(e)(1)** — Transmission Security

ISO-27001-2022

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

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1548	Abuse Elevation Control Mechanism	Privilege-Escalation
T1203	Exploitation for Client Execution	Execution
T1059	Command and Scripting Interpreter	Execution
T1190	Exploit Public-Facing Application	Initial-Access
T1078	Valid Accounts	Defense-Evasion
T1562	Impair Defenses	Defense-Evasion
T1134	Access Token Manipulation	Defense-Evasion
T1565	Data Manipulation	Impact

Sources

Source	URL	Tier
Microsoft Security Blog	https://www.microsoft.com/en-us/security/blog/2026/05/14/defense-in...	T1
	https://govciomedia.com/what-new-guidance-says-for-securing-agentic...	T3

Source	URL	Tier
	https://cyberscoop.com/cisa-nsa-five-eyes-guidance-secure-deploymen...	T3
	https://www.paloaltonetworks.com/blog/network-security/integrated-d...	T3
RCE vulnerabilities in AI agent frameworks Microsoft Security Blog	https://www.microsoft.com/en-us/security/blog/2026/05/07/prompts-be...	T1

DISCLAIMER

This intelligence report is produced by Tech Jacks Solutions Security Command Center (SCC) for informational purposes only. It does not constitute professional security advice, legal counsel, or an incident response engagement. The information herein is derived from publicly available sources and AI-assisted analysis; while every effort is made to ensure accuracy, Tech Jacks Solutions makes no warranties regarding completeness or timeliness. Organizations should conduct their own validation and consult qualified security professionals before taking action based on this report. Tech Jacks Solutions is not liable for any damages resulting from the use of this information.

Generated 2026-05-14 13:50 UTC by TJS Security Command Center