INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-06-09 19:29 UTC

Microsoft Publishes AI Incident Response Playbook; AI System Forensics Emerges as Enterprise IR Capability Gap

SECURITY ANALYSIS | MEDIUM | CVSS 5.0

SCC Item ID	SCC-STY-2026-0181
Type	Security Analysis
Severity	MEDIUM
CVSS Base Score	5.0
Affected Products	Microsoft 365 Copilot, Azure AI Services, Microsoft Sentinel, Microsoft Purview, Microsoft Defender
Published	2026-06-09T17:35:06+00:00
Discovery Source	Rss:T1 Threatintel

Executive Summary

Microsoft published a formal incident response playbook for AI system forensics, acknowledging a documented gap in enterprise capability to reconstruct AI session activity during investigations involving Microsoft 365 Copilot and Azure AI services. The publication arrives as AI-assisted workflows become operationally embedded across enterprises, making the absence of standardized AI forensic methods a compounding liability rather than a future concern. Concurrent movement from Proofpoint on autonomous AI investigation signals that the security industry is treating AI activity reconstruction as an emerging required capability, not an optional enhancement.

Technical Analysis

Microsoft's June 2026 IR playbook targets a specific forensic problem: when an AI system like Microsoft 365 Copilot is involved in a security incident, defenders have lacked the tooling and schema knowledge to reconstruct what the AI did, what data it accessed, and what outputs it generated. The playbook delivers KQL queries, log schema references, and detection logic designed to surface AI interaction telemetry distributed across Microsoft Purview, Defender, and Sentinel, three platforms that do not natively present AI session data in a unified, investigation-ready format.

The core gap is architectural. Enterprise AI systems generate interaction telemetry, but that telemetry is scattered across audit logs, data loss prevention records, and threat detection signals. Without a deliberate forensic methodology, incident responders face the same challenge that plagued cloud forensics a decade ago:

the data exists, but the investigative workflow to assemble it does not. Microsoft's playbook represents the first structured attempt to close that gap within the Microsoft security ecosystem.

The MITRE ATT&CK techniques associated with this story span a broad range, including T1213 (Data from Information Repositories), T1119 (Automated Collection), T1530 (Data from Cloud Storage), and T1041 (Exfiltration Over C2 Channel), which reflect the realistic abuse scenarios for compromised or manipulated AI systems: unauthorized data collection, exfiltration of outputs, and lateral movement through AI-mediated access to enterprise repositories. T1562.001 (Impair Defenses: Disable or Modify Tools) is also listed, consistent with scenarios where adversaries attempt to suppress AI audit logging before or during an attack.

The Miasma campaign reference in the source data provides corroborating threat context. Microsoft's June 2 publication on that campaign documented a supply-chain-adjacent credential-stealing operation targeting npm packages in Red Hat environments, with preinstall persistence mechanisms. While the Miasma campaign is distinct from the IR playbook publication, its appearance in the same source cluster suggests Microsoft is responding to active threat patterns that include credential theft and supply chain compromise techniques that could plausibly intersect with AI system abuse, particularly through T1195.001 (Compromise Software Dependencies and Development Tools), T1552 (Unsecured Credentials), and T1078 (Valid Accounts).

The CWE identifiers in the source data (CWE-312 cleartext storage, CWE-522 insufficiently protected credentials, CWE-798 hardcoded credentials, CWE-494 download without integrity check, CWE-506 embedded malicious code) are best understood as a risk taxonomy for AI system architectures generally, not as confirmed vulnerabilities tied to a specific disclosed flaw. The absence of a CVE identifier and the editorially estimated CVSS score of 0.0 Medium support this interpretation. Defenders should treat these CWEs as a design review checklist for AI system integrations, not as indicators of active exploitation in a named product.

Proofpoint's concurrent announcement of a source-agnostic autonomous investigations platform, while commercially motivated, adds industry signal: two major security vendors are now explicitly positioning AI activity reconstruction as a compliance and IR capability gap requiring dedicated tooling. That convergence matters more than either announcement in isolation.

Action Checklist

1. Step 1: Assess exposure, determine if your organization has deployed Microsoft 365 Copilot, Azure AI Services, or any AI-assisted workflow that generates telemetry ingested by Microsoft Sentinel, Purview, or Defender; document which business units have active Copilot licenses
2. Step 2: Review controls, verify that AI interaction audit logging is enabled in Microsoft Purview (per NIST AU-2, Event Logging, and CIS 8.2, Collect Audit Logs); confirm that Purview Data Loss Prevention policies cover Copilot-generated outputs and that Sentinel is ingesting Purview audit signals
3. Step 3: Obtain and operationalize Microsoft's IR playbook, retrieve the KQL queries and schema references from the Microsoft Security Blog publication dated 2026-06-09; validate that your Sentinel workspace has the required tables populated before an incident forces you to discover gaps under pressure
4. Step 4: Audit AI system credential and configuration hygiene, review Azure AI service principal credentials and API keys for hardcoded or cleartext storage patterns (CWE-798, CWE-312) consistent with the risk taxonomy in the source data; apply NIST AC-6, Least Privilege, to AI service identities and restrict data repository access to minimum required scopes

5. Step 5: Update threat model, incorporate AI session manipulation and AI-mediated data collection (MITRE T1213, T1119, T1530) into your threat register; map these to your Copilot and Azure AI deployment architecture and assign detection coverage ownership
6. Step 6: Communicate findings, brief security leadership and business unit owners with active Copilot deployments on the forensic gap, the availability of Microsoft's playbook, and the timeline for implementation; frame as a logging readiness issue, not a vendor vulnerability
7. Step 7: Monitor developments, track Microsoft Security Blog and Microsoft Purview changelog for updates to the IR playbook, schema changes, or new KQL coverage; monitor for follow-on guidance from CISA on AI system logging requirements as the regulatory posture on AI governance matures

IR / Forensic Enrichment

Triage Priority	STANDARD
Escalation Criteria	Escalate to urgent if investigation of a confirmed or suspected data exfiltration incident involving Microsoft 365 Copilot or Azure AI services reveals that Purview AI interaction audit logs are absent, incomplete, or not ingested into Sentinel, as this constitutes an active forensic evidence gap during a live incident and may trigger breach notification obligations if Copilot-grounded data included PII, PHI, or regulated financial data.
Recovery Notes	Post-implementation, validate logging completeness by conducting a supervised test session in Microsoft 365 Copilot with a known prompt and grounded document, then confirm the corresponding 'CopilotInteraction' record appears in the Purview unified audit log within 15 minutes and is visible in Sentinel within the ingestion SLA of the Purview connector. Monitor Purview audit log volume for AI interaction event types weekly for the first 30 days after enabling logging to detect gaps caused by connector failures, schema changes, or licensing changes that silently disable audit capture. Retain AI interaction audit logs for a minimum of 90 days in Purview and 1 year in Sentinel or equivalent log storage, consistent with NIST AU-11 (Audit Record Retention), to support retroactive investigation of AI-mediated data collection that may not be detected until weeks after occurrence.

Forensic Artifacts	Microsoft Purview Unified Audit Log — 'CopilotInteraction' record type: captures Copilot prompt text, response content, grounded document references (SharePoint URLs, OneDrive file IDs), user identity (UPN), session timestamp, and client IP; primary artifact for reconstructing what data a user or compromised identity caused Copilot to access and synthesize | Microsoft Entra ID Sign-In Logs — filter on application display name 'Microsoft 365 Copilot' and 'Azure Cognitive Services': records authentication events for AI service access including IP address, conditional access policy evaluation results, and MFA status; relevant for determining whether a credential compromise (CWE-798 hardcoded key) was used to authenticate AI service sessions | Azure Activity Log — filter on 'Microsoft.CognitiveServices/accounts/listKeys/action' and 'Microsoft.CognitiveServices/accounts/write' operations: captures who accessed Azure AI API keys, when, and from which IP; directly surfaces unauthorized key access consistent with the CWE-798/CWE-312 credential exposure risk identified in Step 4 | Microsoft Sentinel — 'OfficeActivity' and 'MicrosoftGraphActivityLogs' tables: Graph API calls to '/me/drive', '/sites', or '/users/{id}/messages' made by Azure AI service principals or Copilot service identities map to MITRE T1213 (Data from Information Repositories) and T1530 (Data from Cloud Storage); query for service principal object IDs rather than user UPNs to surface non-human AI-mediated data access | Azure Diagnostic Logs for Cognitive Services — 'AzureDiagnostics' table with ResourceType 'COGNITIVESERVICES/ACCOUNTS': records API call volume, caller IP, operation name, and response codes per Azure AI endpoint; anomalous spikes in call volume or calls from unexpected IP ranges indicate potential T1119 (Automated Collection) behavior where an adversary is using compromised AI service credentials to bulk-query grounded data sources

Per-Action IR Details

Step 1: Assess exposure — determine if your organization has deployed Microsoft 365 Copilot, Azure AI Services, or any AI-assisted workflow that generates telemetry ingested by Microsoft Sentinel, Purview, or Defender; document which business units have active Copilot licenses

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: establishing IR capability and asset visibility before an incident occurs

Controls: NIST AC-2 (Account Management) — enumerate Copilot-licensed accounts and Azure AI service principal identities across the tenant, CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory) — Copilot licenses and Azure AI service endpoints must appear in the asset inventory with owning business unit annotated, CIS 2.1 (Establish and Maintain a Software Inventory) — Microsoft 365 Copilot and Azure AI service deployments require inventory entries with version and scope of data access documented

Compensating: Run 'Get-MgSubscribedSku | Where-Object { $_.SkuPartNumber -like '*Copilot*' }' via Microsoft Graph PowerShell (free, no license required beyond tenant access) to enumerate active Copilot SKUs and assigned users; for Azure AI, run 'az cognitiveservices account list --output table' via Azure CLI to inventory all AI service endpoints and their resource groups — both commands producible by a 2-person team in under 30 minutes.

Evidence: Before scoping, capture a point-in-time snapshot of: Microsoft Entra ID audit logs showing Copilot license assignment events (Microsoft Purview Audit search, Activity: 'Add member to group' filtered to Copilot-linked groups); Azure Activity Log entries for 'Microsoft.CognitiveServices/accounts/write' resource operations; and the output of 'az cognitiveservices account list' to establish a baseline of AI service endpoints prior to any remediation actions that might alter the deployment state.

Step 2: Review controls — verify that AI interaction audit logging is enabled in Microsoft Purview (per NIST AU-2, Event Logging, and CIS 8.2, Collect Audit Logs); confirm that Purview Data Loss Prevention policies cover Copilot-generated outputs and that Sentinel is ingesting Purview audit signals

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: ensuring logging infrastructure is configured to support future detection and forensic reconstruction of AI session activity

Controls: NIST AU-2 (Event Logging) — AI interaction events (Copilot prompts, responses, grounded data references) must be defined as auditable event types and enabled in Microsoft Purview Audit (Standard or Premium), NIST AU-3 (Content Of Audit Records) — Purview audit records for Copilot sessions must capture: what prompt was submitted, when, by which identity, against which data sources, and what response content was generated, NIST AU-9 (Protection Of Audit Information) — Purview audit logs forwarded to Sentinel must be protected from modification; confirm the Sentinel Log Analytics workspace has immutability or retention lock configured, CIS 8.2 (Collect Audit Logs) — validate that Microsoft Purview Audit is enabled at tenant level and that the 'CopilotInteraction' and 'AipSensitivityLabelAction' event types are actively logging to the unified audit log

Compensating: Without Sentinel, use the Microsoft Purview Compliance Portal (free with M365 E3/E5) to run manual Audit Log searches scoped to 'Copilot' activities; export results to CSV and parse with PowerShell for anomalies — schedule this as a weekly cron-equivalent task via Windows Task Scheduler calling 'Search-UnifiedAuditLog -RecordType CopilotInteraction' via Exchange Online PowerShell module (no additional cost).

Evidence: Capture before remediating logging gaps: current Purview Audit configuration state via 'Get-AdminAuditLogConfig' (Exchange Online PowerShell); Sentinel workspace data connector status for 'Microsoft Purview' connector showing last ingestion timestamp; a sample export of 'CopilotInteraction' records from the unified audit log for the past 30 days to establish what was and was not captured prior to any configuration changes — this export is your forensic baseline for any retroactive investigation.

Step 3: Obtain and operationalize Microsoft's IR playbook — retrieve the KQL queries and schema references from the Microsoft Security Blog publication dated 2026-06-09; validate that your Sentinel workspace has the required tables populated before an incident forces you to discover gaps under pressure

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: pre-positioning IR tooling, playbooks, and validated detection queries so they are operable before an incident is declared

Controls: NIST AU-6 (Audit Record Review, Analysis, And Reporting) — KQL queries from the Microsoft playbook operationalize this control by enabling structured, repeatable review of Copilot and Azure AI session audit records in Sentinel, NIST AU-7 (Audit Record Reduction And Report Generation) — Sentinel workbooks and KQL-based analytics rules derived from the Microsoft playbook satisfy the requirement for on-demand report generation from AI session audit data

Compensating: Teams without Sentinel can replicate KQL logic using PowerShell against exported Purview CSV audit logs — translate each KQL WHERE clause to a PowerShell 'Where-Object' filter; store the translated scripts in a shared IR folder and validate them against a known-good test dataset (e.g., a supervised Copilot session with documented prompt/response) to confirm the query returns expected results before relying on them in an incident.

Evidence: Before deploying playbook queries, document which Sentinel tables currently contain data by running 'search * | summarize count() by Type | sort by count_ desc' in Sentinel Log Analytics — capture this table inventory as a screenshot or exported CSV; this establishes which schema references from the Microsoft playbook (e.g., 'OfficeActivity', 'AuditLogs', 'MicrosoftGraphActivityLogs') are populated versus empty, identifying forensic blind spots before an incident exposes them.

Step 4: Audit AI system credential and configuration hygiene — review Azure AI service principal credentials and API keys for hardcoded or cleartext storage patterns (CWE-798, CWE-312) consistent with the risk taxonomy in the source data; apply NIST AC-6, Least Privilege, to AI service identities and restrict data repository access to minimum required scopes

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: hardening AI service identities and credentials to reduce attack surface before exploitation occurs

Controls: NIST AC-6 (Least Privilege) — Azure AI service principals and API keys must be scoped to minimum required permissions; Copilot service accounts must not hold tenant-wide read access to SharePoint or Exchange beyond declared business need, NIST AC-2 (Account Management) — Azure AI service principals are accounts and

require formal lifecycle management: creation justification, periodic review, and revocation when no longer needed, NIST AC-3 (Access Enforcement) — enforce that Azure AI Cognitive Services API keys cannot access data repositories (SharePoint, OneDrive, Exchange) beyond the data scopes explicitly authorized in the service's access control policy, CIS 4.7 (Manage Default Accounts on Enterprise Assets and Software) — audit Azure AI deployments for default or auto-provisioned service principal configurations that retain excessive permissions from initial deployment

Compensating: Run 'az keyvault secret list' and 'az cognitiveservices account keys list' via Azure CLI to audit which secrets are stored in Key Vault versus potentially hardcoded in app configurations; use Microsoft Defender for Cloud's free tier 'Recommendations' blade filtered to 'Cognitive Services' to surface CWE-798/CWE-312 patterns without requiring a paid CSPM; for code repositories, run 'git log -p | grep -i "api.key\|AZURE_KEY\|subscription.key"' to detect cleartext credential commits in repos hosting AI integration code.

Evidence: Before remediating credentials, capture: Azure Entra ID audit log entries for service principal credential creation and last-used timestamps ('Get-MgAuditLogSignIn -Filter "appDisplayName eq \'Azure Cognitive Services\'"); Azure Activity Log entries for 'Microsoft.CognitiveServices/accounts/listKeys/action' showing who has accessed API keys and when; and a snapshot of current Azure role assignments for each AI service principal via 'az role assignment list --assignee --output table' — this documents the pre-remediation permission state for post-incident review.

Step 5: Update threat model — incorporate AI session manipulation and AI-mediated data collection (MITRE T1213, T1119, T1530) into your threat register; map these to your Copilot and Azure AI deployment architecture and assign detection coverage ownership

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: updating threat models and detection ownership to account for AI-specific attack vectors prior to an incident

Controls: NIST AU-6 (Audit Record Review, Analysis, And Reporting) — detection coverage for T1213 (Data from Information Repositories via Copilot grounding), T1119 (Automated Collection via AI-mediated bulk queries), and T1530 (Data from Cloud Storage via Azure AI data connectors) must be assigned to named owners with defined review frequency, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — AI-specific MITRE technique coverage gaps are vulnerability findings in the detection layer and must be tracked in the vulnerability management process with remediation timelines

Compensating: Document the threat model update in a shared spreadsheet mapping each MITRE technique (T1213, T1119, T1530) to specific Copilot or Azure AI data flows in your environment, the current detection method (or 'none'), the log source that would surface it, and the assigned owner; use MITRE ATT&CK Navigator (free, browser-based) to annotate your coverage layer and export as JSON for version-controlled tracking — a 2-person team can complete this mapping in a half-day workshop.

Evidence: Before finalizing the threat model, pull historical Purview unified audit log records for 'CopilotInteraction' events with high data volume (filter on response content length or number of grounded document references) to identify whether T1119-style bulk AI-assisted collection has already occurred in your environment; also query Azure Monitor logs for anomalous Cognitive Services API call volumes using 'AzureDiagnostics | where ResourceType == "COGNITIVESERVICES/ACCOUNTS" | summarize count() by bin(TimeGenerated, 1h), CallerIPAddress' to establish a baseline against which future anomalies can be measured.

Step 6: Communicate findings — brief security leadership and business unit owners with active Copilot deployments on the forensic gap, the availability of Microsoft's playbook, and the timeline for implementation; frame as a logging readiness issue, not a vendor vulnerability

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: lessons learned and capability gap communication to leadership to drive program improvement before the next incident

Controls: NIST AC-1 (Policy And Procedures) — the forensic gap in AI session logging constitutes a policy gap; this communication step initiates the process of updating IR policy to formally cover AI-assisted workflows and define logging requirements

Compensating: Prepare a one-page gap summary (no special tooling required) structured as: current state (what Purview logs and what it does not), risk statement (what an investigator cannot reconstruct without the missing logs),

required action (enable AI audit logging in Purview, validate Sentinel ingestion), and timeline; use the Microsoft Security Blog post dated 2026-06-09 as the external authority reference to anchor the urgency — business unit owners respond better to a vendor-published playbook than to internal security assertions alone.

Evidence: Before the brief, compile a concrete evidence package showing the gap: run a Purview Audit search for 'CopilotInteraction' events for the past 90 days and document the result count (zero or sparse results confirm the logging gap is real); include the Sentinel table inventory from Step 3 showing unpopulated AI-relevant tables; this evidence transforms an abstract capability gap into a measurable, demonstrable finding that leadership can act on with appropriate urgency.

Step 7: Monitor developments — track Microsoft Security Blog and Microsoft Purview changelog for updates to the IR playbook, schema changes, or new KQL coverage; monitor for follow-on guidance from CISA on AI system logging requirements as the regulatory posture on AI governance matures

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: continuous improvement through intelligence integration and monitoring of evolving guidance relevant to identified capability gaps

Controls: NIST AU-13 (Monitoring For Information Disclosure) — monitoring Microsoft Security Blog and CISA AI guidance channels operationalizes this control by tracking external sources for schema changes or new disclosure patterns that would affect Copilot forensic coverage, CIS 7.2 (Establish and Maintain a Remediation Process) — tracking playbook updates and new KQL coverage from Microsoft feeds directly into the remediation process for AI logging gaps identified in Steps 2 and 3

Compensating: Configure an RSS feed reader (free: Feedly, NewsBlur, or raw RSS via curl in a cron job) to ingest 'https://www.microsoft.com/en-us/security/blog/feed/' and 'https://www.cisa.gov/news-events/cybersecurity-advisories/rss.xml'; set keyword alerts for 'Copilot', 'Purview audit', 'AI forensics', and 'KQL schema' — route matches to a shared IR inbox and review weekly; no SIEM or paid tooling required.

Evidence: Maintain a versioned changelog document tracking: date of each Microsoft Purview schema update affecting 'CopilotInteraction' record fields, date and content of each Microsoft Security Blog AI IR playbook revision, and any CISA advisory or binding operational directive referencing AI system logging — this changelog serves as the evidentiary record that your organization maintained awareness of evolving guidance, which is directly relevant to regulatory inquiries about AI governance due diligence.

Detection Guidance

Primary log sources for AI forensic investigation within the Microsoft ecosystem are Microsoft Purview Audit (unified audit log), Microsoft Defender for Cloud Apps, and Microsoft Sentinel. Per NIST AU-2 and CIS 8.2, confirm audit logging is enabled and collecting AI interaction events before building detection logic against gaps.

Key telemetry to validate and hunt against:

- Purview Unified Audit Log: Look for CopilotInteraction and AIActivityEvent record types. Gaps in expected interaction frequency from active Copilot users may indicate log suppression (T1562.001). Bulk data access events within Copilot sessions warrant review against T1213 and T1119.

- Microsoft Sentinel: Apply the KQL queries from Microsoft's 2026-06-09 IR playbook to reconstruct AI session timelines. Prioritize queries that correlate Copilot interaction records with SharePoint and OneDrive access events to identify anomalous data collection patterns consistent with T1530.

- Defender for Cloud Apps: Review AI app governance policies and alert on Copilot sessions that access sensitive labeled content outside normal user behavior baselines. Correlate with Purview sensitivity labels to identify potential exfiltration via AI-generated outputs (T1041).

- Credential and identity signals: Given the Miasma campaign context and T1078, T1552 presence in the technique set, audit Azure AD sign-in logs for AI service principal authentications from unexpected IPs or outside normal operational windows. Apply D3-LAM (Local Account Monitoring) principles to AI service identities. Review for hardcoded credentials in AI integration code using static analysis (CWE-798, CWE-522).

- Supply chain and integrity signals: For organizations using npm or similar package ecosystems in AI toolchain development, validate package integrity per CWE-494 patterns identified in the Miasma campaign. Apply D3-FMBV (File Magic Byte Verification) and D3-SFA (System File Analysis) to AI pipeline dependencies.

Hunting hypothesis: A user account with valid credentials (T1078) accesses Microsoft 365 Copilot, directs it to enumerate sensitive SharePoint repositories (T1213), and the output is exfiltrated through normal Copilot response channels without triggering DLP because the data leaves via an AI-generated document summary rather than a direct file download. Hunt for Copilot sessions that access high-sensitivity labeled content followed by external sharing events within the same session window.

Indicators of Compromise

Type	Value	Context	Confidence
URL	`Pending — refer to Microsoft Security Blog (2026-06-02) for published Miasma campaign indicators`	Microsoft published IOCs including npm package names, file hashes, and persistence mechanism details associated with the Miasma credential-stealing campaign targeting Red Hat npm environments; the actual indicator values are not reproduced in the provided source data	LOW

Framework Mappings

MITRE-ATTACK

- **T1566** — Phishing
- **T1098** — Account Manipulation
- **T1553.002** — Code Signing
- **T1552.001** — Credentials In Files
- **T1190** — Exploit Public-Facing Application
- **T1530** — Data from Cloud Storage
- **T1552** — Unsecured Credentials
- **T1119** — Automated Collection
- **T1078** — Valid Accounts
- **T1213** — Data from Information Repositories
- **T1195.001** — Compromise Software Dependencies and Development Tools
- **T1041** — Exfiltration Over C2 Channel
- **T1562.001** — Disable or Modify Tools
- **T1059.007** — JavaScript

NIST-800-53R5

- **AT-2** — Literacy Training and Awareness
- **CA-7** — Continuous Monitoring
- **SC-7** — Boundary Protection
- **SI-3** — Malicious Code Protection
- **SI-4** — System Monitoring
- **SI-8** — Spam Protection
- **CA-8** — Penetration Testing
- **RA-5** — Vulnerability Monitoring and Scanning
- **SI-2** — Flaw Remediation
- **SI-7** — Software, Firmware, and Information Integrity
- **AC-6** — Least Privilege
- **AC-2** — Account Management
- **IA-2** — Identification and Authentication (Organizational Users)
- **IA-5** — Authenticator Management
- **CM-3** — Configuration Change Control

OWASP-TOP10-2021

- **A08:2021** — Software and Data Integrity Failures
- **A04:2021** — Insecure Design
- **A07:2021** — Identification and Authentication Failures

CIS-V8

- **2.5** — Allowlist Authorized Software
- **2.6** — Allowlist Authorized Libraries
- **5.2** — Use Unique Passwords
- **16.10** — Apply Secure Design Principles in Application Architectures
- **6.3** — Require MFA for Externally-Exposed Applications
- **14.2** — Train Workforce Members to Recognize Social Engineering Attacks
- **8.2** — Collect Audit Logs

HIPAA-SECURITY

- **164.308(a)(5)(ii)(D)** — Password Management
- **164.312(d)** — Person or Entity Authentication
- **164.312(e)(1)** — Transmission Security

ISO-27001-2022

- **A.8.28** — Secure coding
- **A.5.23** — Information security for use of cloud services

SOC2-TSC

- **CC6.1** — Logical access security software, infrastructure, and architectures

NIST-CSF-2

- **DE.CM-01** — Networks and network services are monitored

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1566	Phishing	Initial-Access
T1098	Account Manipulation	Persistence
T1553.002	Code Signing	Defense-Evasion
T1552.001	Credentials In Files	Credential-Access
T1190	Exploit Public-Facing Application	Initial-Access
T1530	Data from Cloud Storage	Collection
T1552	Unsecured Credentials	Credential-Access
T1119	Automated Collection	Collection
T1078	Valid Accounts	Defense-Evasion
T1213	Data from Information Repositories	Collection
T1195.001	Compromise Software Dependencies and Development Tools	Initial-Access
T1041	Exfiltration Over C2 Channel	Exfiltration
T1562.001	Disable or Modify Tools	Defense-Evasion
T1059.007	JavaScript	Execution

Sources

Source	URL	Tier
Microsoft Security Blog	https://www.microsoft.com/en-us/security/blog/2026/06/09/reconstruc...	T1
	https://www.microsoft.com/en-us/security/blog/2026/06/09/reconstruc...	T1
	https://www.proofpoint.com/us/newsroom/press-releases/proofpoint-es...	T3
	https://www.microsoft.com/en-us/security/blog/2026/06/02/preinstall...	T1

Source	URL	Tier
Microsoft Security Copilot	https://www.microsoft.com/en-us/security/business/ai-machine-learni...	T1

Generated 2026-06-09 19:29 UTC by TJS Security Command Center