INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-07-08 07:04 UTC

WriteOut: Writer Enterprise AI Platform Flaw Enabled Cross-Tenant Session Hijacking via Agent Preview Links

SECURITY ANALYSIS | CRITICAL | CVSS 9.5

SCC Item ID	SCC-STY-2026-0332
Type	Security Analysis
Severity	CRITICAL
CVSS Base Score	9.5
Affected Products	Writer (enterprise generative AI platform), specific version not publicly disclosed; flaw existed in agent preview proxy layer; now patched
Published	2026-07-07T09:27:09
Discovery Source	Rss

Executive Summary

A critical vulnerability in Writer's enterprise AI platform, dubbed 'WriteOut' by Sand Security researchers, allowed unauthenticated attackers to hijack authenticated user sessions across organizational boundaries by distributing a malicious agent preview link. No prior access to the victim's tenant was required, and successful exploitation granted full account takeover. The flaw exposes a systemic risk in multi-tenant SaaS architecture: preview and sandbox subsystems frequently lack the isolation guarantees of the primary application, and enterprise AI platforms introduce new attack surface that security programs have not yet fully mapped.

Technical Analysis

Sand Security researchers identified a session hijacking vulnerability in Writer's agent preview proxy layer, the component responsible for rendering AI agent outputs before deployment. According to reporting by The Hacker News, the proxy failed to enforce tenant isolation, meaning a preview link generated in one tenant could be weaponized against a user in a completely separate organization.

The attack chain required no prior foothold in the victim's environment. An attacker crafts a malicious preview link and delivers it to a target user, via spearphishing or any other delivery mechanism (MITRE T1566.002). When the authenticated victim clicks the link, the preview proxy processes the request without validating tenant boundaries, exposing the victim's session token (CWE-384, CWE-668). The attacker captures that token and achieves full account takeover (T1539, T1550.004) with the victim's privileges inside Writer's platform.

The CVSS base score of 9.5 reflects the attack profile: network-accessible, low complexity, no authentication required, with high impact on confidentiality and integrity. The assigned CWEs, session fixation/hijacking (CWE-384), improper restriction of communication channels (CWE-923), exposure of resource to wrong sphere (CWE-668), and protection mechanism failure (CWE-693), collectively describe a failure to extend core security assumptions into a derived subsystem.

The broader implication for security teams is architectural. Enterprise AI platforms expose new internal components, agent sandboxes, preview renderers, output proxies, that inherit trust from the primary application without necessarily inheriting its access controls. This is the same class of isolation failure that has historically plagued cloud storage preview handlers, email sandbox detonation environments, and browser extension APIs. As enterprises accelerate AI platform adoption, these peripheral subsystems represent an expanding and under-audited attack surface. The vulnerability is now patched; no CVE identifier has been publicly assigned at this time.

Action Checklist

1. Step 1: Assess exposure, determine whether your organization uses Writer's enterprise AI platform; if so, contact your Writer account representative or support team to confirm that the patched version is deployed in your tenant and obtain written confirmation of the patch date.
2. Step 2: Review session token controls, audit whether your enterprise AI platform vendors enforce tenant isolation at the preview, sandbox, and proxy layer, not only at the primary application layer; request architecture documentation or a vendor security questionnaire response covering these subsystems (aligns with NIST AC-4: Information Flow Enforcement).
3. Step 3: Enforce and verify MFA across all enterprise AI platform accounts, session token theft loses impact when token reuse alone cannot complete authentication; verify MFA enrollment for all Writer and analogous platform accounts (aligns with CIS 6.3: Require MFA for Externally-Exposed Applications; D3-MFA: Multi-factor Authentication).
4. Step 4: Audit access and account privileges, review which users hold elevated permissions inside Writer and peer platforms; apply least privilege to limit the blast radius of any future session compromise (aligns with NIST AC-6: Least Privilege; CIS 5.4: Restrict Administrator Privileges; D3-UAP: User Account Permissions).
5. Step 5: Update your threat model, add 'AI platform preview/sandbox subsystem isolation failure' as an explicit threat scenario in your SaaS risk register; extend vendor security review processes to require explicit attestation of tenant isolation for all subsystems, not only primary application endpoints (aligns with NIST AC-20: Use of External Systems).
6. Step 6: Brief leadership, frame this for business leadership as a supply-chain trust risk: enterprise AI tools integrated into core workflows may carry session-level access risks that bypass traditional perimeter controls; pair with a review of AI platform procurement and security review standards.
7. Step 7: Monitor developments, track for CVE assignment, additional Sand Security technical disclosures, and any follow-on regulatory guidance regarding AI platform security from CISA or equivalent authorities.

IR / Forensic Enrichment

Triage Priority	URGENT
Escalation Criteria	Escalate immediately to CISO and legal counsel if IdP or Writer audit logs reveal any user account received and redeemed a Writer agent preview link from an unrecognized or external sender prior to patch confirmation, as this constitutes probable session compromise and may trigger breach notification obligations under applicable data protection regulations (e.g., GDPR 72-hour notification, CCPA, or sector-specific requirements) given Writer's integration with enterprise content and potentially sensitive organizational data.
Recovery Notes	After confirming Writer has deployed the patched version in your tenant (Step 1) and MFA enforcement is validated (Step 3), force-expire all active Writer sessions via your IdP SSO session policy or Writer admin console to ensure no pre-patch hijacked tokens remain valid. Monitor Writer tenant audit logs and IdP session logs for anomalous access patterns for a minimum of 30 days post-patch — specifically watch for session originations from unexpected geographic locations or IP ranges inconsistent with your workforce, which would indicate a hijacked session established before containment. Revalidate the privilege audit (Step 4) at the 30-day mark to confirm no unauthorized role escalations occurred during the exposure window.
Forensic Artifacts	IdP session logs (Okta System Log, Azure AD Sign-In Logs) filtered for Writer app assignments — specifically events for 'user.session.start', 'app.oauth2.token.grant', and 'user.authentication.sso' within the 90-day window before patch confirmation, to detect cross-tenant session redemptions originating from unexpected tenants or unauthenticated sources | Writer tenant admin audit log — events covering preview link generation, agent preview access, workspace access, document access, and role changes; the WriteOut flaw operated at the agent preview proxy layer, so preview link creation and redemption events are the highest-value artifacts for establishing whether exploitation occurred | Browser forensic artifacts on endpoints of users who may have clicked a Writer agent preview link from an external or unexpected sender — Chrome/Edge Cookies SQLite DB (path: %LOCALAPPDATA%\Google\Chrome\User Data\Default\Network\Cookies) for writer.com and writer preview subdomain cookies, and browser history for writer.com preview URL patterns | Network proxy or DNS logs for connections to Writer preview and agent subdomain endpoints (e.g., preview.writer.com or equivalent agent proxy hostnames) — specifically requests that originated without a preceding authenticated session establishment, which would indicate exploitation of the unauthenticated preview proxy flaw | Email and collaboration platform logs (Exchange message trace, Slack audit logs, Teams audit logs) for inbound messages containing Writer agent preview links originating from external or unrecognized senders — the WriteOut attack vector required distributing a malicious preview link, making delivery channel logs a primary source for identifying targeted users and establishing the attack timeline

Per-Action IR Details

Step 1: Assess exposure — determine whether your organization uses Writer's enterprise AI platform; if so, confirm with your Writer account representative that the patched version is deployed in your tenant and obtain written confirmation of the patch date.

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis: scope assessment and asset triage to determine whether the organization is affected by a disclosed vulnerability

Controls: CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory), CIS 2.1 (Establish and Maintain a Software Inventory), CIS 2.2 (Ensure Authorized Software is Currently Supported)

Compensating: Query your SaaS application inventory (or review SSO/IdP application catalog in Okta, Azure AD, or Google Workspace) to enumerate all Writer tenant registrations and associated user accounts. If no formal SaaS inventory exists, search IT procurement records and expense reports for Writer subscriptions. Document tenant IDs and workspace URLs for every confirmed instance.

Evidence: Before any session revocation or credential rotation, capture current Writer session token artifacts: export active SSO session records from your IdP (e.g., Okta System Log — event type 'user.session.start' and 'app.oauth2.token.grant'), and preserve any browser session cookies for Writer domains (app.writer.com, preview subdomains) from potentially affected endpoints using browser forensic tools (e.g., BrowsingHistoryView or manual SQLite extraction from Chrome's Cookies DB at %LOCALAPPDATA%\Google\Chrome\User Data\Default\Network\Cookies). These are volatile and will be destroyed upon session revocation in Step 3.

Step 2: Review session token controls — audit whether your enterprise AI platform vendors enforce tenant isolation at the preview, sandbox, and proxy layer, not only at the primary application layer; request architecture documentation or a vendor security questionnaire response covering these subsystems (aligns with NIST AC-4: Information Flow Enforcement).

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: establishing vendor security review requirements and validating third-party controls before or during incident assessment

Controls: NIST AC-4 (Information Flow Enforcement), NIST AC-20 (Use Of External Systems)

Compensating: Issue a written security questionnaire to your Writer account representative specifically asking: (1) Does the agent preview proxy layer enforce per-tenant session token binding? (2) Are preview subdomain sessions cryptographically scoped to the originating tenant? (3) What audit logging exists for preview link generation and redemption events? Document responses with timestamps. If Writer cannot provide architecture documentation, treat tenant isolation as unverified and escalate to your vendor risk management process.

Evidence: No live state alteration occurs in this step; volatile capture is not required. However, preserve existing vendor contract terms, SLAs, and any prior security questionnaire responses for the Writer platform as baseline documentation — these establish pre-incident knowledge and are relevant to post-incident vendor accountability and regulatory disclosure timelines.

Step 3: Enforce and verify MFA across all enterprise AI platform accounts — session token theft loses impact when token reuse alone cannot complete authentication; verify MFA enrollment for all Writer and analogous platform accounts (aligns with CIS 6.3: Require MFA for Externally-Exposed Applications; D3-MFA: Multi-factor Authentication).

NIST Phase: Containment

Reference: NIST 800-61r3 §3.3 — Containment Strategy: reducing the impact of an active or potential session hijacking threat by enforcing authentication controls that limit token reuse value

Controls: CIS 6.3 (Require MFA for Externally-Exposed Applications), CIS 6.5 (Require MFA for Administrative Access), NIST AC-7 (Unsuccessful Logon Attempts)

Compensating: Pull an MFA enrollment report from your IdP (Okta: Admin → Reports → User MFA Enrollment; Azure AD: Azure Portal → Users → Authentication Methods Activity). Filter for all users with Writer app assignments and flag any accounts where MFA factor is 'not enrolled' or factor type is SMS-only (phishable). For accounts without IdP-brokered SSO to Writer, contact Writer support to export per-tenant MFA enrollment status. Prioritize enforcement for accounts with admin or workspace-owner roles identified in Step 4.

Evidence: CRITICAL — volatile capture required before session revocation: Before forcing session termination or MFA re-enrollment (which invalidates existing tokens), export the full active session list from your IdP for Writer-assigned users (Okta System Log query: eventType eq 'user.session.start' AND target.displayName eq 'Writer', last 30 days). Additionally, if any user received and clicked a Writer agent preview link from an unrecognized source, capture browser process memory and network connection state using Task Manager → Create Dump File on the browser process, and run 'netstat -ano' or 'Get-NetTCPConnection' to record active connections to Writer preview domains before killing those sessions.

Step 4: Audit access and account privileges — review which users hold elevated permissions inside Writer and peer platforms; apply least privilege to limit the blast radius of any future session compromise (aligns with NIST AC-6: Least Privilege; CIS 5.4: Restrict Administrator Privileges; D3-UAP: User Account Permissions).

NIST Phase: Containment

Reference: NIST 800-61r3 §3.3 — Containment Strategy: limiting blast radius by reducing the privilege footprint accessible via a hijacked cross-tenant session

Controls: NIST AC-6 (Least Privilege), NIST AC-2 (Account Management), CIS 5.4 (Restrict Administrator Privileges to Dedicated Administrator Accounts), CIS 5.1 (Establish and Maintain an Inventory of Accounts)

Compensating: Request a full user and role export from your Writer tenant admin console (Settings → Users → Export). Identify all accounts assigned 'Admin', 'Workspace Owner', or equivalent elevated roles. Cross-reference against your IdP group memberships to detect any role assignments that exceed what was approved during provisioning. For each over-privileged account, downgrade to standard user role via Writer Admin → Users → Edit Role, and document the change with justification. Repeat this review for any peer generative AI platforms (e.g., Jasper, Notion AI) integrated into the same SSO federation.

Evidence: Before modifying any account permissions in Writer (which may alter audit trail state), export the current Writer admin audit log covering the past 90 days — specifically events for role assignment changes, workspace creation, and API key generation. If Writer provides a SCIM or SIEM integration, pull provisioning event logs showing when elevated roles were granted. These records establish the pre-remediation privilege baseline and are essential for determining whether an attacker who successfully hijacked a session via the WriteOut flaw gained access to an admin-level account.

Step 5: Update your threat model — add 'AI platform preview/sandbox subsystem isolation failure' as an explicit threat scenario in your SaaS risk register; extend vendor security review processes to require explicit attestation of tenant isolation for all subsystems, not only primary application endpoints (aligns with NIST AC-20: Use of External Systems).

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: lessons learned and policy updates that improve detection and prevention of the same vulnerability class in future vendor assessments

Controls: NIST AC-20 (Use Of External Systems), CIS 7.1 (Establish and Maintain a Vulnerability Management Process), CIS 7.2 (Establish and Maintain a Remediation Process)

Compensating: Add a mandatory line item to your vendor security review questionnaire template (even if maintained in a spreadsheet): 'Does the platform enforce tenant-boundary isolation for all subsystems including preview, sandbox, proxy, and developer/test environments? Provide architecture evidence.' Tag Writer and all generative AI SaaS vendors in your risk register with the new threat scenario 'SaaS Preview Subsystem Isolation Failure — Cross-Tenant Session Hijacking' and set residual risk review cadence to quarterly until vendors provide architecture attestation.

Evidence: No volatile evidence capture required for this documentation step. However, retain the Sand Security WriteOut research disclosure, your Writer vendor patch confirmation obtained in Step 1, and all communications with Writer's account team as supporting artifacts for the updated risk register entry. These records substantiate the threat scenario addition and provide regulatory audit trail if a breach notification obligation is later assessed.

Step 6: Brief leadership — frame this for business leadership as a supply-chain trust risk: enterprise AI tools integrated into core workflows may carry session-level access risks that bypass traditional perimeter controls; pair with a review of AI platform procurement and security review standards.

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: communicating incident impact and systemic risk to leadership to drive policy and procurement improvements

Controls: NIST AC-1 (Policy And Procedures)

Compensating: Prepare a one-page executive brief using the following specific framing for WriteOut: an unauthenticated external attacker could distribute a single malicious Writer agent preview link and achieve full account

takeover of any Writer user who clicked it — with no prior access to your tenant required. Quantify blast radius using the account inventory from Step 4 (number of users, including any with admin roles, document access scope, and any sensitive data processed through Writer workflows). Pair with a request to add generative AI SaaS platforms to the formal vendor security review gate in the procurement process.

Evidence: No volatile forensic capture required. Attach the privilege audit output from Step 4 and the patch confirmation from Step 1 as supporting exhibits to the leadership brief to substantiate both the risk and the remediation status.

Step 7: Monitor developments — track for CVE assignment, additional Sand Security technical disclosures, and any follow-on regulatory guidance regarding AI platform security from CISA or equivalent authorities.

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: integrating updated threat intelligence and regulatory guidance into ongoing detection and response improvement

Controls: CIS 7.1 (Establish and Maintain a Vulnerability Management Process)

Compensating: Create a monitoring task (even a calendar reminder or RSS feed subscription) for the following: (1) NVD/MITRE CVE database for any new CVE assigned to the Writer WriteOut flaw (search 'Writer AI' or 'writer.com'); (2) Sand Security's blog or disclosure repository for technical indicators, PoC details, or additional affected endpoints that could generate IOCs for retrospective log review; (3) CISA Known Exploited Vulnerabilities catalog and CISA AI security advisories for regulatory guidance. If a CVE is assigned and a CVSS or EPSS score becomes available, re-triage priority against your patching SLA policy.

Evidence: No volatile forensic capture required for this monitoring step. If Sand Security releases a technical disclosure with specific request patterns (e.g., malformed preview link URL parameters, specific HTTP headers, or token formats used in the WriteOut attack), immediately re-examine Writer tenant access logs and IdP session logs for the 90-day window preceding the patch confirmation date obtained in Step 1 to identify any retrospective exploitation attempts.

Detection Guidance

Session anomaly detection is the primary hunt surface for this attack pattern. Review authentication and session logs for the following behavioral indicators, sourced from the attack chain described in Sand Security's research:

- Unexpected session token usage: look for the same session token appearing from two distinct IP addresses or geographic locations within a short time window (indicative of token theft and reuse, T1550.004).

- Abnormal access patterns within Writer or peer AI platforms: access from unfamiliar user agents, ASNs, or geolocations immediately following a link-click event should trigger review.

- Preview or agent endpoint access without a corresponding authenticated session initiation: if your platform logs are granular enough, cross-reference preview proxy requests against session establishment events; orphaned preview requests warrant investigation.

- Spearphishing delivery vectors: review email gateway and proxy logs for Writer preview link distribution to internal users from external or untrusted senders (T1566.002).

Log sources to prioritize: IdP/SSO authentication logs, SaaS CASB telemetry for Writer, email gateway logs for Writer-domain preview link delivery, and any available Writer audit logs (request audit log access from your Writer account team if not currently ingested).

For threat hunting, build a hypothesis around T1539 (Steal Web Session Cookie) and T1550.004 (Web Session Cookie reuse): hunt for session tokens active on multiple concurrent source IPs, or sessions that initiated via a link click rather than a standard login flow.

NIST AU-6 (Audit Record Review, Analysis, and Reporting) and AU-12 (Audit Record Generation) provide the control framework for ensuring these log sources are collected and reviewed. CIS 8.2 (Collect Audit Logs) establishes the baseline expectation that SaaS platform audit logs are ingested into centralized logging.

No specific IOC values (hashes, IPs, domains) were present in the provided source material. The Sand Security blog post at https://www.sandsecurity.ai/blog/writeout-writer-ai-cross-tenant may contain additional technical indicators; consult that source directly for values if conducting active threat hunting.

Framework Mappings

MITRE-ATTACK

- **T1185** — Browser Session Hijacking
- **T1071.001** — Web Protocols
- **T1550.004** — Web Session Cookie
- **T1059** — Command and Scripting Interpreter
- **T1190** — Exploit Public-Facing Application
- **T1566.002** — Spearphishing Link
- **T1539** — Steal Web Session Cookie

NIST-800-53R5

- **CM-7** — Least Functionality
- **SI-3** — Malicious Code Protection
- **SI-4** — System Monitoring
- **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
- **AT-2** — Literacy Training and Awareness
- **SI-8** — Spam Protection

CIS-V8

- **7.3** — Perform Automated Operating System Patch Management
- **7.4** — Perform Automated Application Patch Management

ISO-27001-2022

- **A.8.8** — Management of technical vulnerabilities
- **A.5.23** — Information security for use of cloud services

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1185	Browser Session Hijacking	Collection
T1071.001	Web Protocols	Command-And-Control
T1550.004	Web Session Cookie	Defense-Evasion
T1059	Command and Scripting Interpreter	Execution
T1190	Exploit Public-Facing Application	Initial-Access
T1566.002	Spearphishing Link	Initial-Access
T1539	Steal Web Session Cookie	Credential-Access

Sources

Source	URL	Tier
Security News	https://thehackernews.com/2026/07/writer-ai-flaw-could-let-agent-pr...	T2
World-class enterprises trust WRITER	https://writer.com/trust/	T3
WriteOut Flaw in Writer AI Enabled Cross-Tenant Session Hijacking	https://www.mallory.ai/stories/019f3cea-cb32-76c7-8f6a-9fc26e967179	T3
WriteOut: Abusing the Sandbox for a Critical Cross-Tenant ...	https://www.sandsecurity.ai/blog/writeout-writer-ai-cross-tenant	T3

Generated 2026-07-08 07:04 UTC by TJS Security Command Center