INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-05-01 18:56 UTC

AI Blind Spots: Why Enterprise AI Inventories Are Wrong and What That Means for Security Teams

GOVERNANCE | MEDIUM | CVSS 5.0

SCC Item ID	SCC-GOV-2026-0027
Type	Governance
Severity	MEDIUM
CVSS Base Score	5.0
Affected Products	Enterprise AI deployments (cross-platform); CrowdStrike Falcon platform, Shadow AI Visibility Service, Falcon AIDR, CrowdStrike AI Security Services
Discovery Source	Rss:T1 Threatintel

Executive Summary

According to CrowdStrike professional services engagements, organizations frequently undercount their AI deployments; one documented case found a customer tracking 150 AI agents actually operating over 500. Untracked AI deployments inherit existing user permissions without authorization review, creating unsanctioned data access pathways across endpoint, SaaS, and cloud environments. Without an accurate AI asset inventory, risk assessment and access controls cannot function, exposing organizations to data leakage, compliance gaps, and access control failures they cannot currently see.

Technical Analysis

The core failure is an absent or inaccurate AI asset inventory, creating downstream breakdowns in access governance and data protection. Shadow AI deployments, including MCP servers, IDE extensions, and unsanctioned SaaS AI tools, inherit the permissions of the user or service account that installed or invoked them (CWE-284), without explicit authorization review. Prompt-layer interactions with sensitive data introduce exfiltration vectors across endpoint, SaaS, and cloud surfaces (CWE-200). There is no single CVE; the risk is architectural and governance-based. Relevant MITRE ATT&CK techniques include T1078 (Valid Accounts, AI agents inheriting user permissions), T1048 (Exfiltration Over Alternative Protocol, prompt-based data exfiltration), T1530 (Data from Cloud Storage, AI agent data access), T1059 (Command and Scripting Interpreter, agent-executed commands), and T1195 (Supply Chain Compromise, MCP server and IDE extension vectors). CrowdStrike has released the Shadow AI Visibility Service, expanded Falcon AIDR detection capabilities, and unified data protection controls targeting this gap.

Action Checklist

1. Step 1: Inventory, run discovery against your environment to enumerate all AI agents, SaaS AI integrations, IDE extensions, MCP servers, and browser-based AI tools. Compare discovered inventory against self-reported registers to identify gaps. CrowdStrike's Shadow AI Visibility Service is one option; similar discovery can be performed via network flow analysis, SaaS access reviews, and endpoint agent telemetry.
2. Step 2: Detection, review identity and access management logs for service accounts or user accounts with broad data permissions that are associated with AI tool activity; look for anomalous data access patterns from API tokens issued to AI integrations; flag MCP server registrations and IDE plugin installations in endpoint telemetry.
3. Step 3: Eradication, for each discovered AI deployment, conduct an explicit authorization review; revoke or scope down permissions that were inherited without review; disable or block unsanctioned AI tools that cannot be brought under policy; treat unreviewed AI agents as unauthorized access pathways until cleared.
4. Step 4: Recovery, validate that the revised AI asset inventory is accurate and maintained through an automated discovery process, not manual tracking; confirm that access controls for approved AI agents are least-privilege and documented; monitor for re-emergence of shadow deployments via ongoing SaaS and endpoint telemetry.
5. Step 5: Post-Incident, the root control gap here is the absence of an AI asset management process equivalent to traditional IT asset management; implement a formal AI governance policy that requires authorization review before deployment, periodic re-discovery audits, and data classification alignment for any AI tool touching sensitive data.

IR / Forensic Enrichment

Triage Priority	STANDARD
Escalation Criteria	Escalate to urgent and engage legal counsel if Step 3 authorization review reveals that any unreviewed AI agent had documented access to PII, PHI, payment card data, or data subject to contractual confidentiality obligations, as this may trigger breach notification obligations under GDPR, HIPAA, CCPA, or applicable state law — the unauthorized access pathway existed for an undetermined period prior to discovery.
Recovery Notes	Post-remediation, maintain automated AI asset discovery as a continuous process rather than a one-time exercise, since the CrowdStrike engagement data shows shadow AI deployments re-emerge as users install new tools; run re-discovery against SaaS OAuth grants and endpoint telemetry at minimum monthly for the first quarter post-remediation. Verify that the approved AI agent list is reviewed against vendor security advisories on the same cadence as traditional software, as AI agent frameworks (LangChain, AutoGen, MCP) release security-relevant updates that affect the risk profile of approved deployments. Watch specifically for re-emergence of IDE AI extensions on developer workstations and new MCP server registrations in endpoint telemetry, as these were the highest-volume shadow deployment vectors identified in the CrowdStrike engagement findings.

Forensic Artifacts	SaaS OAuth grant logs (Microsoft 365 Unified Audit Log 'ConsentToApplication' events; Google Workspace Token Audit log) — these record every instance where a user authorized an AI tool to access organizational data, including the granted permission scopes and authorizing account, which directly evidence the unauthorized access pathways created by unreviewed AI deployments | Entra ID or Okta service principal and app role assignment records — export via 'Get-MgServicePrincipalAppRoleAssignment' PowerShell cmdlet; these records show which AI agent service accounts were granted which data permissions without authorization review, and at what privilege level | Endpoint software and extension inventory snapshots from MDM (Intune/Jamf) or osquery ('SELECT name, path FROM browser_extensions' and installed apps queries) — establishes the pre-remediation AI tool footprint on developer and knowledge worker endpoints, including MCP server binaries and IDE AI plugins not visible in SaaS logs | Network DNS query logs or proxy logs for AI provider API domains (openai.com, anthropic.com, api.cohere.ai, generativelanguage.googleapis.com, *.bedrock.amazonaws.com) — these reveal AI tool usage by endpoint and volume, identifying high-data-access AI integrations that may not appear in OAuth logs because they use API keys embedded in applications rather than user-delegated OAuth grants | Cloud storage and collaboration access logs (SharePoint/OneDrive access logs in M365 Compliance Center; Google Drive audit log) filtered on service principal or API token principals matching AI tool identities — these are the primary evidence source for what sensitive data unreviewed AI agents actually accessed, and are essential for any breach notification assessment

Per-Action IR Details

Step 1: Inventory — run discovery against your environment to enumerate all AI agents, SaaS AI integrations, IDE extensions, MCP servers, and browser-based AI tools; do not rely on self-reported registers. CrowdStrike's Shadow AI Visibility Service is one option; similar discovery can be performed via network flow analysis, SaaS access reviews, and endpoint agent telemetry.

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis: establishing scope of adverse event by enumerating all affected assets before analysis can be accurate

Controls: NIST SI-4 (System Monitoring) — monitor for unapproved software and service registrations across endpoint and network layers, NIST CM-8 (System Component Inventory) — maintain a current, accurate inventory of system components including AI agents and SaaS integrations, CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory) — AI agents, MCP servers, and IDE extensions are enterprise assets and must be enumerated, CIS 2.1 (Establish and Maintain a Software Inventory) — IDE AI plugins (GitHub Copilot, Cursor, Codeium) and browser extensions constitute installed software requiring inventory coverage

Compensating: For teams without CrowdStrike Falcon: (1) Deploy osquery enterprise-wide and run 'SELECT name, path, permissions FROM browser_extensions' plus 'SELECT name, path FROM apps' to surface AI browser extensions and desktop agents. (2) Export OAuth app authorizations from Google Workspace Admin (Admin Console > Security > API Controls > App Access Control) or Microsoft Entra ID (Enterprise Applications > All Applications, filter by 'User consent') to enumerate SaaS AI integrations granted data access. (3) Use Zeek or Wireshark with a BPF filter for DNS queries resolving to known AI provider domains (openai.com, anthropic.com, api.cohere.ai, huggingface.co, *.amazonaws.com/bedrock) to detect network-layer AI traffic not represented in self-reported inventories.

Evidence: Before running discovery tooling that may alter state: (1) Export a point-in-time snapshot of all OAuth tokens and API keys currently authorized in your IdP (Entra ID sign-in logs, Google Workspace token audit log) — these will disappear if revoked during Step 3 before they are recorded. (2) Capture current SaaS audit logs showing which user accounts have granted AI tool permissions (Salesforce Setup Audit Trail, Microsoft 365 Unified Audit Log event 'ConsentToApplication'). (3) Pull endpoint software inventory from existing MDM (Intune, Jamf) or AV console before any remediation action changes the installed state — this establishes the baseline the 3x undercounting finding applies to.

Step 2: Detection — review identity and access management logs for service accounts or user accounts with broad data permissions that are associated with AI tool activity; look for anomalous data access patterns from API tokens issued to AI integrations; flag MCP server registrations and IDE plugin installations in endpoint telemetry.

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis: correlating IAM log sources and endpoint telemetry to characterize the scope of unauthorized access pathways created by untracked AI deployments

Controls: NIST AU-6 (Audit Record Review, Analysis, and Reporting) — review IAM and API gateway logs specifically for service accounts and OAuth tokens associated with AI tool activity, NIST AU-2 (Event Logging) — ensure logging is enabled for OAuth consent events, API token issuance, and SaaS data-access events that AI integrations would generate, NIST IR-5 (Incident Monitoring) — track and document each AI deployment identified as an unauthorized access pathway, NIST SI-4 (System Monitoring) — monitor endpoint telemetry for MCP server process registrations and IDE extension installations, CIS 5.1 (Establish and Maintain an Inventory of Accounts) — enumerate all service accounts and API tokens associated with AI tools, as these represent undocumented access principals

Compensating: Without a SIEM: (1) In Microsoft 365 Unified Audit Log, run a Search-UnifiedAuditLog PowerShell query filtering on RecordType 'AzureActiveDirectory' and Operations 'Add service principal', 'Consent to application', 'Add app role assignment to service principal' — these are the exact event types generated when a user grants an AI SaaS tool access to M365 data. (2) For Google Workspace, pull the Token Audit log (Admin Console > Reports > Audit > Token) filtered to third-party apps with Drive or Gmail scopes — this surfaces AI tools like Notion AI, Grammarly, or ChatGPT plugins with broad data grants. (3) On endpoints, use Sysmon Event ID 1 (Process Creation) to detect MCP server processes (e.g., 'mcp-server-*.exe', 'npx @modelcontextprotocol/*') and Sysmon Event ID 11 (File Create) for IDE extension directories under '%APPDATA%\Code\extensions' (VS Code) or '%APPDATA%\JetBrains'.

Evidence: Before any access revocation: (1) Export the full OAuth grant list for each AI tool identified in Step 1, including granted scopes and the authorizing user account — this documents the inherited permission set that existed without authorization review. (2) Capture API gateway or cloud trail logs showing data access volume by AI-associated tokens: AWS CloudTrail 'GetObject' events by principal matching AI service account names; Azure Monitor logs filtering on service principal activity for AI app registrations. (3) Record Sysmon or EDR telemetry showing MCP server process trees and parent-child relationships to IDE processes, establishing which developer workstations are running unapproved MCP configurations.

Step 3: Eradication — for each discovered AI deployment, conduct an explicit authorization review; revoke or scope down permissions that were inherited without review; disable or block unsanctioned AI tools that cannot be brought under policy; treat unreviewed AI agents as unauthorized access pathways until cleared.

NIST Phase: Eradication

Reference: NIST 800-61r3 §3.4 — Eradication: removing unauthorized access pathways by revoking inherited permissions and disabling unsanctioned AI agents, analogous to removing malicious persistence mechanisms

Controls: NIST IR-4 (Incident Handling) — execute eradication actions as part of the formal incident handling process, documenting each authorization decision, NIST AC-2 (Account Management) — disable or scope-down service accounts and API tokens associated with AI tools that failed authorization review, NIST AC-6 (Least Privilege) — revoke overly broad permissions inherited by AI agents without authorization review; re-scope to minimum required data access, NIST CM-7 (Least Functionality) — disable or block unsanctioned AI tools at the endpoint and network layer; prohibit execution of unapproved MCP server processes, CIS 6.2 (Establish an Access Revoking Process) — apply a formal access revocation process to all AI tool service accounts and OAuth tokens that were granted without review, CIS 5.4 (Restrict Administrator Privileges to Dedicated Administrator Accounts) — verify no AI agent service accounts hold administrative or elevated permissions inherited through role assignment

Compensating: Without an enterprise PAM or IGA tool: (1) In Entra ID, use 'Remove-MgServicePrincipalAppRoleAssignment' and 'Revoke-MgUserSignInSession' PowerShell cmdlets to revoke specific app role assignments for AI tools; document each revocation with ticket reference before executing. (2) In Google Workspace, use the Admin SDK Directory API or Admin Console Token page to revoke OAuth tokens per-user per-application — automate with a Python script using google-auth and googleapiclient libraries iterating over the token audit export from Step 2. (3) Block unsanctioned AI tool domains at the DNS layer using Pi-hole or your existing DNS

forwarder with a blocklist for API endpoints of unapproved tools (api.openai.com, claude.ai/api, generativelanguage.googleapis.com) — this prevents re-authorization even if endpoint controls are bypassed.

Evidence: Before revoking any permission: (1) Screenshot or log-export the full permission scope of each AI agent being revoked — this is your forensic record of the unauthorized access pathway that existed. For Entra ID, run 'Get-MgServicePrincipalAppRoleAssignment -ServicePrincipalId ' and export to CSV. (2) Preserve IAM audit log entries showing the original consent grant event, including the authorizing user, timestamp, and granted scopes — needed if a breach notification assessment is required later. (3) Document whether any AI agent had access to data classified as PII, PHI, or regulated financial data, as this determines whether the unauthorized access constitutes a reportable incident under HIPAA, GDPR, or CCPA.

Step 4: Recovery — validate that the revised AI asset inventory is accurate and maintained through an automated discovery process, not manual tracking; confirm that access controls for approved AI agents are least-privilege and documented; monitor for re-emergence of shadow deployments via ongoing SaaS and endpoint telemetry.

NIST Phase: Recovery

Reference: NIST 800-61r3 §3.5 — Recovery: restoring a known-good state by validating inventory accuracy and access control integrity, then establishing monitoring to detect regression to the shadow-AI condition

Controls: NIST IR-4 (Incident Handling) — confirm recovery actions are complete and the environment is returned to a controlled state with documented AI access controls, NIST SI-7 (Software, Firmware, and Information Integrity) — validate that only authorized AI agents with documented permission scopes remain active post-remediation, NIST AU-12 (Audit Record Generation) — ensure ongoing logging is configured to capture new AI tool OAuth grants, API token issuance, and MCP server registrations going forward, NIST CM-8 (System Component Inventory) — confirm the AI asset inventory is now maintained through automated discovery and reflects the post-remediation state, CIS 7.2 (Establish and Maintain a Remediation Process) — verify that the remediation process used in Step 3 is documented and repeatable for future shadow AI discoveries, CIS 8.2 (Collect Audit Logs) — validate that audit logging now covers AI-specific event types: OAuth consent events, API token creation, and SaaS application installations

Compensating: For ongoing monitoring without a SIEM: (1) Schedule a weekly osquery query via cron or Task Scheduler to re-run the browser extension and installed app inventory from Step 1, diff against the approved baseline, and email the delta to the security team — this is your shadow AI re-emergence detector. (2) Configure a weekly export of the Microsoft 365 or Google Workspace OAuth app authorization report and compare against the approved AI tool list; new entries not on the approved list trigger a Step 2 investigation. (3) Use a Sigma rule converted to a native query for your log platform targeting process creation events for known MCP server binary names and VS Code extension installation paths to catch re-emergence on developer endpoints.

Evidence: During recovery validation: (1) Re-run the osquery software inventory and OAuth export from Step 1 post-remediation and retain the diff as evidence that unauthorized AI deployments were removed — this is your remediation effectiveness record. (2) Confirm in Entra ID or Google Workspace that no service principal or OAuth grant exists for previously revoked AI tools — export the post-revocation app authorization state as a timestamped snapshot. (3) Verify in endpoint telemetry (Sysmon, EDR, or MDM) that blocked MCP server processes and IDE extensions are no longer executing or present on developer workstations flagged in Step 2.

Step 5: Post-Incident — the root control gap here is the absence of an AI asset management process equivalent to traditional IT asset management; implement a formal AI governance policy that requires authorization review before deployment, periodic re-discovery audits, and data classification alignment for any AI tool touching sensitive data.

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: using lessons learned from this engagement to close the structural control gap — absence of AI-specific asset management — that enabled 3x undercounting of AI deployments

Controls: NIST IR-8 (Incident Response Plan) — update the IR plan to include AI asset discovery as a standing component of preparation and periodic review cycles, NIST IR-2 (Incident Response Training) — train security and IT staff to recognize shadow AI deployment patterns: OAuth grants to unfamiliar SaaS apps, MCP server processes, IDE

extension mass-installs, NIST RA-3 (Risk Assessment) — conduct a formal risk assessment for each approved AI agent class, incorporating data classification alignment for tools with access to sensitive data, NIST SI-2 (Flaw Remediation) — treat absence of authorization review for AI deployments as a flaw in the access control process requiring documented remediation with target completion dates, NIST CM-8 (System Component Inventory) — formalize AI agents as a tracked component class in the enterprise asset inventory, equivalent to servers and endpoints, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — extend vulnerability management scope to include AI agents: track versions, vendor advisories, and model update events as vulnerability-relevant changes, CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory) — formally add AI agents, MCP servers, and SaaS AI integrations as asset classes with required inventory fields: owner, data access scope, authorization date, and review frequency

Compensating: For organizations without a GRC platform: (1) Create an AI Tool Authorization Register as a shared spreadsheet or wiki page with mandatory fields: tool name, owner, data classification of accessible data, OAuth scopes granted, authorization approver, approval date, and next review date — enforce via a lightweight approval workflow using your existing ticketing system (Jira, ServiceNow, or even GitHub Issues). (2) Write a one-page AI Acceptable Use Policy addendum that explicitly prohibits connecting AI tools to systems containing PII, PHI, or regulated data without a completed authorization entry in the register — distribute via your existing security awareness platform. (3) Schedule quarterly re-discovery runs using the osquery and OAuth export process from Step 1 as a formal audit procedure, with results reviewed by the security team and delta items opened as remediation tickets.

Evidence: For post-incident documentation and lessons learned: (1) Produce a final count comparing the pre-engagement self-reported AI inventory against the discovery-verified count — this quantifies the 3x+ undercounting gap in your specific environment and is the primary metric for the lessons-learned report. (2) Retain a copy of all authorization review decisions made in Step 3 (approved, scoped-down, or revoked) as the baseline for the first formal AI asset register. (3) Document which data classification tiers were accessible to unreviewed AI agents — if PII, PHI, or payment card data was reachable, this finding may require legal review for breach notification assessment under GDPR Article 33, HIPAA 45 CFR §164.400, or applicable state law.

Detection Guidance

There is no single IOC or CVE signature for this governance gap. Detection is behavioral and inventory-based. Query endpoint telemetry for AI-related process names, browser extensions, and IDE plugins across the fleet. Review SaaS access logs for OAuth grants issued to AI applications, flag any grant with broad data scopes (mail.read, files.readwrite, calendars.read) that was not explicitly approved. Audit identity provider logs for service accounts or API tokens associated with AI tool activity. In cloud environments, review IAM policies for principals that represent AI agents or automation; flag any with permissions exceeding what the documented use case requires. Falcon AIDR (CrowdStrike) provides detection coverage for AI-specific data leakage behaviors if the Falcon agent is deployed. For organizations without purpose-built tooling, network flow analysis showing data egress to AI provider endpoints (OpenAI, Anthropic, Cohere, Azure OpenAI, etc.) from unexpected sources is a practical starting point.

Framework Mappings

MITRE-ATTACK

- **T1048** — Exfiltration Over Alternative Protocol
- **T1195** — Supply Chain Compromise
- **T1530** — Data from Cloud Storage
- **T1078** — Valid Accounts
- **T1059** — Command and Scripting Interpreter

- **T1526** — Cloud Service Discovery
- **T1552** — Unsecured Credentials
- **T1087** — Account Discovery

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
- **AC-2** — Account Management
- **AC-6** — Least Privilege
- **IA-2** — Identification and Authentication (Organizational Users)
- **IA-5** — Authenticator Management
- **CM-7** — Least Functionality
- **SI-3** — Malicious Code Protection
- **SI-4** — System Monitoring
- **AC-3** — Access Enforcement
- **SC-28** — Protection of Information at Rest

OWASP-TOP10-2021

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

HIPAA-SECURITY

- **164.312(a)(1)** — Access Control
- **164.308(a)(6)(ii)** — Response and Reporting

CIS-V8

- **6.1** — Establish an Access Granting Process
- **6.2** — Establish an Access Revoking Process
- **8.2** — Collect Audit Logs

SOC2-TSC

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

NIST-CSF-2

- **RS.CO-03** — Recovery activities and progress communicated
- **DE.CM-01** — Networks and network services are monitored

ISO-27001-2022

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

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1048	Exfiltration Over Alternative Protocol	Exfiltration
T1195	Supply Chain Compromise	Initial-Access
T1530	Data from Cloud Storage	Collection
T1078	Valid Accounts	Defense-Evasion
T1059	Command and Scripting Interpreter	Execution
T1526	Cloud Service Discovery	Discovery
T1552	Unsecured Credentials	Credential-Access
T1087	Account Discovery	Discovery

Sources

Source	URL	Tier
Blog	https://www.crowdstrike.com/en-us/blog/crowdstrike-shadow-AI-visibi...	T3
	https://www.crowdstrike.com/en-us/blog/crowdstrike-secures-growing-...	T3
	https://www.crowdstrike.com/en-us/blog/crowdstrike-stops-genai-data...	T3
	https://siliconangle.com/2026/03/23/crowdstrike-targets-ai-security...	T3
CrowdStrike Shadow AI Visibility Service | Reduce AI Footprint Risk	https://www.crowdstrike.com/en-us/services/ai-security-services/sha...	T3

Generated 2026-05-01 18:56 UTC by TJS Security Command Center