INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-06-17 13:59 UTC

# AI Agents Need Identity Too: CrowdStrike's Continuous Authorization Model Targets Non-Human Privilege Risk

**SECURITY ANALYSIS** | **MEDIUM** | CVSS 5.0

| | |
|---|---|
| **SCC Item ID** | SCC-STY-2026-0218 |
| **Type** | Security Analysis |
| **Severity** | MEDIUM |
| **CVSS Base Score** | 5.0 |
| **Affected Products** | CrowdStrike Falcon Next-Gen Identity Security, CrowdStrike Falcon AI Detection and Response (AIDR), Falcon Zero Trust Access (ZTA), AWS cloud infrastructure |
| **Discovery Source** | Rss:T1 Threatintel |

## Executive Summary

CrowdStrike has introduced Continuous Identity for AI Agents, a real-time authorization framework that evaluates every action taken by autonomous AI agents against current identity, device posture, and business context rather than relying on static, pre-granted permissions. The announcement signals a structural shift in enterprise identity security: as AI agents proliferate across cloud and SaaS environments, the authenticate-once model that governs human access creates exploitable standing privilege windows that threat actors can abuse through token theft, account manipulation, and application access token hijacking. Organizations deploying agentic AI without corresponding identity governance controls are accumulating non-human privilege debt that existing IAM programs are not designed to detect or remediate.

## Technical Analysis

The core problem CrowdStrike is addressing is architectural, not product-specific. Enterprise AI agents, like service accounts before them, tend to accumulate privileges at provisioning time and retain them indefinitely. Unlike service accounts, AI agents act autonomously across dynamic workflows, making the standing-privilege model exponentially more dangerous: a compromised agent token or misconfigured OAuth scope does not wait for a human to log in before acting.

CrowdStrike's Continuous Identity for AI Agents applies zero-trust principles to non-human identities by evaluating authorization at the moment of each agent action. The framework is built on two open standards: SPIFFE (Secure Production Identity Framework for Everyone), which provides cryptographically verifiable

workload identities, and the OpenID Foundation's Shared Signals Framework (SSF), which enables real-time posture signals to flow between systems and inform access decisions continuously rather than at session initiation.

The weakness patterns the framework targets map directly to well-documented MITRE ATT&CK techniques. T1078 (Valid Accounts) and T1078.004 (Cloud Accounts) describe adversaries using legitimate credentials, including over-provisioned agent tokens, to blend into normal traffic. T1550.001 (Application Access Token) covers the abuse of OAuth tokens and API keys, which are the primary credential type for non-human identities. T1528 (Steal Application Access Token) and T1098 (Account Manipulation) complete the picture: an adversary who can harvest an agent's long-lived token gains persistent, policy-invisible access to whatever that agent was authorized to reach, including data exfiltration via cloud storage access (T1530).

The CWE patterns are equally telling. CWE-250 (Execution with Unnecessary Privileges) and CWE-269 (Improper Privilege Management) reflect the over-provisioning that occurs when agent permissions are set to 'whatever the workflow might need' rather than scoped to each discrete action. CWE-284 (Improper Access Control) and CWE-522 (Insufficiently Protected Credentials) describe the downstream consequences when those permissions are neither audited nor rotated.

The solution extends these controls to AWS cloud infrastructure and provides unified governance across service accounts, API keys, and OAuth tokens, the three non-human identity types most commonly abused in cloud-native breaches. The reference architecture is vendor-backed, built on open standards, and compatible with the Falcon platform's existing Zero Trust Access and AI Detection and Response capabilities, giving security teams a path to integrate non-human identity monitoring into their existing SIEM and SOAR workflows.

## Action Checklist

1. Step 1: Inventory non-human identities, audit all AI agents, service accounts, API keys, and OAuth tokens across your environment, including AWS and SaaS integrations; map which actions each is authorized to perform (CIS 1.1: Establish and Maintain Detailed Enterprise Asset Inventory; CIS 5.1: Establish and Maintain an Inventory of Accounts)
2. Step 2: Identify standing privilege exposure, flag any agent or service account that holds persistent, broadly scoped permissions rather than time-limited or action-scoped authorizations; treat each as a standing privilege risk analogous to a dormant admin account (NIST AC-6: Least Privilege; CIS 5.3: Disable Dormant Accounts)
3. Step 3: Enforce least privilege and credential rotation on non-human identities, scope agent permissions to the minimum required for each discrete action; implement rotation schedules for API keys and OAuth tokens (NIST AC-6: Least Privilege; NIST AC-2: Account Management; D3-CRO: Credential Rotation; D3-CH: Credential Hardening)
4. Step 4: Extend MFA and access controls to non-human identity workflows where supported, for human-in-the-loop agent approval gates, require MFA at the decision boundary; for fully autonomous agents, enforce cryptographic workload identity (SPIFFE or equivalent) rather than shared secrets (CIS 6.3: Require MFA for Externally-Exposed Applications; CIS 6.5: Require MFA for Administrative Access; NIST AC-3: Access Enforcement; D3-MFA: Multi-factor Authentication)
5. Step 5: Enable logging and behavioral monitoring for agent actions, ensure audit logs capture agent identity, action type, resource accessed, and authorization context for every agent-initiated event; alert on token use from unexpected source IPs or at unusual times (NIST AU-2: Event Logging; NIST AU-3: Content of Audit Records; CIS 8.2: Collect Audit Logs; D3-LAM: Local Account Monitoring; D3-UAP: User

Account Permissions)

6. Step 6: Update your threat model, add T1078.004 (Cloud Accounts), T1550.001 (Application Access Token), and T1528 (Steal Application Access Token) as priority hunt hypotheses for any environment where AI agents or automated workflows hold cloud resource access

7. Step 7: Evaluate vendor roadmap alignment, if your organization uses CrowdStrike Falcon, assess whether Continuous Identity for AI Agents, Falcon AIDR, and Falcon ZTA are in scope for your current deployment; if you use competing platforms, request equivalent non-human identity governance capability documentation from your vendor

## IR / Forensic Enrichment

| | |
|---|---|
| **Triage Priority** | STANDARD |
| **Escalation Criteria** | Escalate to urgent if discovery during the Step 1–2 inventory reveals any AI agent or service account with wildcard cloud permissions (`*:*`) that has authenticated from an IP outside known agent infrastructure within the past 30 days, or if CrowdStrike Falcon AIDR generates a detection on anomalous OAuth token use coinciding with lateral movement indicators in the same AWS account or SaaS tenant. |
| **Recovery Notes** | After completing credential rotation and least-privilege scoping (Steps 3–4), verify all AI agents resume expected operational behavior by reviewing CloudTrail for successful `AssumeRole` and API call patterns matching pre-rotation baselines — unexpected failures indicate over-scoped restriction requiring policy adjustment before declaring recovery complete. Monitor CrowdStrike Falcon Next-Gen Identity Security and CloudTrail for at least 30 days post-remediation for recurrence of broad-scope token use, unauthorized `sts:AssumeRoleWithWebIdentity` calls, or OAuth grants to unrecognized applications. Document all identity scope changes in a change record tied to this remediation effort to support future audit evidence requirements. |
| **Forensic Artifacts** | AWS CloudTrail logs filtered for `AssumeRole`, `AssumeRoleWithWebIdentity`, `GetSessionToken`, and `CreateServiceSpecificCredential` events — these directly evidence AI agent or service account credential use patterns, including source IPs and session durations that would reveal token theft (T1528) or unauthorized application access token reuse (T1550.001) \| AWS IAM Credential Report (`aws iam generate-credential-report`) — captures last-used timestamps and access key age for all service accounts and IAM users, providing forensic evidence of which non-human identities were actively operating and whether any credentials were used after expected rotation windows \| CrowdStrike Falcon Next-Gen Identity Security identity event timeline — contains per-action authorization evaluation records specific to Falcon's Continuous Identity for AI Agents framework, including device posture scores and business context signals at the time of each agent action, which are unavailable from CloudTrail alone \| OAuth token grant history from the identity provider admin console (Microsoft Entra ID audit logs `Add app role assignment to service principal` or Google Workspace Admin SDK Reports API `token` event type) — evidences which applications received delegated permissions from AI agent workflows and the scope granted, directly relevant to T1528 and T1550.001 investigation \| AWS Config configuration change history for IAM roles and policies — records the exact timestamps and content of any permission scope changes on agent roles, establishing whether a policy was modified by an authorized administrator or by an attacker using a compromised agent identity with `iam:PutRolePolicy` or `iam:AttachRolePolicy` permissions |

### Per-Action IR Details

**Step 1: Inventory non-human identities — audit all AI agents, service accounts, API keys, and OAuth tokens across your environment, including AWS and SaaS integrations; map which actions each is authorized to perform (CIS 1.1: Establish and Maintain Detailed Enterprise Asset Inventory; CIS 5.1: Establish and Maintain an Inventory of Accounts)**

**NIST Phase:** Preparation

**Reference:** NIST 800-61r3 §2 — Preparation: establishing visibility into identity assets before an AI agent abuse incident occurs

**Controls:** CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory), CIS 5.1 (Establish and Maintain an Inventory of Accounts), NIST AC-2 (Account Management)

**Compensating:** Run `aws iam list-roles` and `aws iam list-users` filtered for service roles and programmatic-access users; export with `aws iam generate-credential-report` to surface all access keys and last-used timestamps. For SaaS OAuth tokens, pull the Google Workspace or Microsoft 365 token grant list via their respective admin APIs or portal exports. Use a spreadsheet to map each identity to its authorized actions — achievable by a 2-person team in a half-day sprint.

**Evidence:** Before beginning the audit, capture a point-in-time snapshot of active AWS IAM sessions via `aws iam get-account-authorization-details` and export CloudTrail logs covering the prior 90 days of `AssumeRole`, `GetSessionToken`, and `CreateServiceSpecificCredential` events — these volatile usage records establish a pre-audit baseline of which agent identities were actively operating and cannot be reconstructed after credentials are rotated or deleted.

**Step 2: Identify standing privilege exposure — flag any agent or service account that holds persistent, broadly scoped permissions rather than time-limited or action-scoped authorizations; treat each as a standing privilege risk analogous to a dormant admin account (NIST AC-6: Least Privilege; CIS 5.3: Disable Dormant Accounts)**

**NIST Phase:** Detection Analysis

**Reference:** NIST 800-61r3 §3.2 — Detection and Analysis: analyzing identity posture to identify exploitable standing privilege conditions created by the authenticate-once model for AI agents

**Controls:** NIST AC-6 (Least Privilege), CIS 5.3 (Disable Dormant Accounts), CIS 5.4 (Restrict Administrator Privileges to Dedicated Administrator Accounts)

**Compensating:** Use `aws iam get-account-authorization-details` output parsed with `jq` to extract all roles with `*:*` or broad wildcard policies attached to service principals; flag any role whose `RoleLastUsed.LastUsedDate` is null or older than 45 days. For OAuth tokens, export the Microsoft Entra ID or Google Workspace admin console's app permission report and filter on scopes containing `*.ReadWrite.All` or equivalent broad grants. Document each finding in a risk register.

**Evidence:** Before flagging or modifying any standing-privilege identity, capture current AWS CloudTrail `ListAccessKeys` and `GetAccessKeyLastUsed` output for all service accounts, plus a full export of CrowdStrike Falcon's identity inventory if Falcon Next-Gen Identity Security is deployed — this establishes which broadly scoped agent identities were active at the time of analysis and preserves evidence of their authorization scope that would be lost after remediation changes the permission state.

**Step 3: Enforce least privilege and credential rotation on non-human identities — scope agent permissions to the minimum required for each discrete action; implement rotation schedules for API keys and OAuth tokens (NIST AC-6: Least Privilege; NIST AC-2: Account Management; D3-CRO: Credential Rotation; D3-CH: Credential Hardening)**

**NIST Phase:** Containment

**Reference:** NIST 800-61r3 §3.3 — Containment: scoping and rotating credentials to eliminate standing privilege exposure without full eradication, limiting blast radius of any AI agent identity compromise

**Controls:** NIST AC-6 (Least Privilege), NIST AC-2 (Account Management), CIS 5.2 (Use Unique Passwords), CIS 6.2 (Establish an Access Revoking Process)

**Compensating:** Use `aws iam create-policy` to define action-scoped inline policies per agent role, then `aws iam detach-role-policy` to remove the existing broad policy before attaching the scoped replacement. Rotate all IAM access keys using `aws iam create-access-key` followed by `aws iam delete-access-key` on the old key — do this sequentially, not simultaneously, to avoid service disruption. Set a cron job or AWS Lambda scheduled event to enforce 90-day key rotation. For OAuth tokens, revoke and reissue via the identity provider's API.

**Evidence:** Before rotating any credential or detaching any IAM policy, capture a complete `aws iam simulate-principal-policy` output for each agent role to document its current effective permissions, and pull CloudTrail logs for the 24 hours preceding rotation covering `AssumeRole`, `InvokeAPI`, and any `sts:AssumeRoleWithWebIdentity` events — rotating credentials immediately destroys the ability to correlate pre-rotation agent actions to specific permission scopes if an incident investigation is later opened.

**Step 4: Extend MFA and access controls to non-human identity workflows where supported — for human-in-the-loop agent approval gates, require MFA at the decision boundary; for fully autonomous agents, enforce cryptographic workload identity (SPIFFE or equivalent) rather than shared secrets (CIS 6.3: Require MFA for Externally-Exposed Applications; CIS 6.5: Require MFA for Administrative Access; NIST AC-3: Access Enforcement; D3-MFA: Multi-factor Authentication)**

**NIST Phase:** Containment

**Reference:** NIST 800-61r3 §3.3 — Containment: hardening the authentication boundary for AI agent workflows to prevent token reuse and lateral movement via stolen application access tokens

**Controls:** NIST AC-3 (Access Enforcement), NIST AC-17 (Remote Access), CIS 6.3 (Require MFA for Externally-Exposed Applications), CIS 6.5 (Require MFA for Administrative Access)

**Compensating:** For human-in-the-loop gates, enforce MFA via your identity provider's conditional access policy (Entra ID Conditional Access or AWS IAM condition key `aws:MultiFactorAuthPresent: true` on sensitive role assumption). For autonomous agents, deploy SPIFFE/SPIRE (open source, free) to issue short-lived X.509 SVIDs as workload identity documents, replacing long-lived API keys; the SPIRE server can run on a single VM. Document each agent's SVID issuance and renewal interval in your identity inventory.

**Evidence:** Before enforcing new authentication requirements on any externally-exposed agent workflow, capture the current CrowdStrike Falcon ZTA device posture evaluation output and any existing Falcon Next-Gen Identity Security authentication event logs for those agent accounts — if an agent identity was already compromised and operating via a stolen OAuth token (T1528), altering the auth controls without first capturing active session data will destroy evidence of the attacker's current access pattern.

**Step 5: Enable logging and behavioral monitoring for agent actions — ensure audit logs capture agent identity, action type, resource accessed, and authorization context for every agent-initiated event; alert on token use from unexpected source IPs or at unusual times (NIST AU-2: Event Logging; NIST AU-3: Content of Audit Records; CIS 8.2: Collect Audit Logs; D3-LAM: Local Account Monitoring; D3-UAP: User Account Permissions)**

**NIST Phase:** Detection Analysis

**Reference:** NIST 800-61r3 §3.2 — Detection and Analysis: establishing the audit log capability required to detect AI agent identity abuse, including anomalous token use indicative of T1550.001 (Application Access Token) or T1528 (Steal Application Access Token)

**Controls:** NIST AU-2 (Event Logging), NIST AU-3 (Content of Audit Records), NIST AU-6 (Audit Record Review, Analysis, and Reporting), CIS 8.2 (Collect Audit Logs)

**Compensating:** Enable AWS CloudTrail with S3 data event logging and configure an SNS alert for `ConsoleLogin` or `AssumeRole` events sourced from IPs outside your known agent execution environment (Lambda VPC, ECS cluster subnet, or CI/CD runner CIDR). Write a free Sigma rule targeting CloudTrail `eventName: AssumeRoleWithWebIdentity` with `sourceIPAddress` not matching your agent infrastructure CIDRs. Forward logs to an ELK stack (free) or use `jq` scheduled queries against S3-stored CloudTrail JSON for teams without SIEM. Retain logs per NIST AU-11 for a minimum consistent with your retention policy.

**Evidence:** This step establishes forward-looking detection capability; however, before enabling new logging configurations that may overwrite or rotate existing log streams, export and archive current AWS CloudTrail event

history, CrowdStrike Falcon AIDR detection telemetry, and any existing AWS Config configuration change history — these pre-logging-change records may be the only evidence of agent behavior that predates the monitoring implementation and are critical if a retroactive investigation is needed.

**Step 6: Update your threat model — add T1078.004 (Cloud Accounts), T1550.001 (Application Access Token), and T1528 (Steal Application Access Token) as priority hunt hypotheses for any environment where AI agents or automated workflows hold cloud resource access**

**NIST Phase:** Post Incident

**Reference:** NIST 800-61r3 §4 — Post-Incident Activity: updating the threat model and detection hypotheses based on new AI agent identity attack surface understanding, feeding improved detection back into the preparation and detection phases

**Controls:** NIST AU-6 (Audit Record Review, Analysis, and Reporting)

**Compensating:** Document three hunt hypotheses in a shared wiki or runbook: (1) CloudTrail query for `sts:AssumeRoleWithWebIdentity` calls where the federated identity is an AI agent role and the source IP is outside expected agent infrastructure; (2) AWS CloudTrail filter for `eventName: GetSessionToken` followed within 60 seconds by high-volume S3 or DynamoDB API calls from the same session token; (3) Entra ID or Google Workspace audit log query for OAuth token grants to applications not in the approved software inventory (CIS 2.1). Store these as dated, versioned Sigma or KQL rules in version control.

**Evidence:** Before finalizing updated threat model documentation, collect and preserve the current state of CrowdStrike Falcon AIDR's existing detection rule set and any prior Falcon Next-Gen Identity Security alert history related to service account or OAuth activity — these records establish the pre-update detection baseline and are necessary to measure whether the new hunt hypotheses improve true-positive rates in post-implementation review.

**Step 7: Evaluate vendor roadmap alignment — if your organization uses CrowdStrike Falcon, assess whether Continuous Identity for AI Agents, Falcon AIDR, and Falcon ZTA are in scope for your current deployment; if you use competing platforms, request equivalent non-human identity governance capability documentation from your vendor**

**NIST Phase:** Post Incident

**Reference:** NIST 800-61r3 §4 — Post-Incident Activity: capability gap assessment and vendor alignment to ensure the organization's tooling can enforce continuous authorization for AI agents rather than relying on the authenticate-once model that created the standing privilege risk

**Controls:** CIS 7.1 (Establish and Maintain a Vulnerability Management Process), CIS 7.2 (Establish and Maintain a Remediation Process)

**Compensating:** For teams without CrowdStrike Falcon, build a capability gap matrix comparing your current identity platform's non-human identity features against the five core requirements of CrowdStrike's Continuous Identity model: (1) per-action authorization evaluation, (2) device posture integration, (3) business context signals, (4) short-lived credential issuance, (5) behavioral anomaly detection for agents. Score each gap as mitigated, partially mitigated, or unmitigated, and document compensating controls (SPIFFE/SPIRE for workload identity, osquery for device posture, CloudTrail anomaly alerting) for each unmitigated gap. Present the matrix to leadership with a remediation timeline.

**Evidence:** Before concluding the vendor assessment, preserve a point-in-time export of your current CrowdStrike Falcon subscription entitlements and module activation status (available from the Falcon console under Subscription & Usage), and archive any prior Falcon AIDR or Falcon ZTA evaluation reports — these records establish the capability baseline against which post-assessment improvements will be measured and support audit evidence requirements under NIST CA (Assessment, Authorization, and Monitoring) family controls.

## Detection Guidance

Detection for non-human identity abuse requires instrumentation that most organizations have not yet applied to agent and service account activity. Priority log sources: cloud provider identity logs (AWS CloudTrail IAM

events, Azure AD sign-in logs), API gateway access logs, OAuth token issuance and exchange logs, and any agent orchestration platform logs (LangChain, AutoGen, or equivalent).

Behavioral patterns to hunt for:

- Token use anomalies: an agent token authenticating from a source IP, region, or time window inconsistent with its provisioned workflow (maps to T1078, T1078.004)

- Scope creep: an agent or service account accessing resources outside its documented function, particularly cloud storage buckets or secrets managers (maps to T1530, T1550.001)

- Token reuse or replay: the same OAuth or API token appearing in requests that originate from multiple hosts or processes simultaneously (maps to T1528)

- Privilege escalation attempts: agent identities requesting permissions beyond their provisioned scope, even if the request is denied (maps to CWE-269, CWE-250)

- Orphaned tokens: API keys or OAuth tokens associated with decommissioned agents or deprecated workflows that remain active and are still generating authentication events

Policy audit gaps to surface:

- Service accounts and agent identities not included in the quarterly access review cycle

- API keys with no expiration date or rotation policy

- OAuth scopes granted at the broadest available level rather than the minimum required

- Absence of SPIFFE or equivalent cryptographic workload identity for any agent that accesses production systems

For CrowdStrike Falcon customers: Falcon AIDR provides behavioral detection for AI agent activity; Falcon ZTA enforces continuous authorization policy. Review whether agent identities are enrolled in ZTA policy scope, not just human user accounts.

NIST AU-6 (Audit Record Review, Analysis, and Reporting) and AU-12 (Audit Record Generation) apply directly: ensure agent actions generate structured audit records reviewable for anomaly patterns, not just binary allow/deny access logs.

## Framework Mappings

### MITRE-ATTACK

- **T1078** — Valid Accounts
- **T1098** — Account Manipulation
- **T1078.004** — Cloud Accounts
- **T1550.001** — Application Access Token
- **T1528** — Steal Application Access Token
- **T1530** — Data from Cloud Storage

### NIST-800-53R5

- **AC-2** — Account Management
- **AC-6** — Least Privilege

- **IA-2** — Identification and Authentication (Organizational Users)
- **IA-5** — Authenticator Management
- **AC-3** — Access Enforcement

### OWASP-TOP10-2021

- **A01:2021** — Broken Access Control
- **A04:2021** — Insecure Design
- **A07:2021** — Identification and Authentication Failures

### CIS-V8

- **6.1** — Establish an Access Granting Process
- **6.2** — Establish an Access Revoking Process
- **5.2** — Use Unique Passwords
- **5.4** — Restrict Administrator Privileges to Dedicated Administrator Accounts
- **6.8** — Define and Maintain Role-Based Access Control
- **6.3** — Require MFA for Externally-Exposed Applications

### SOC2-TSC

- **CC6.1** — The entity implements logical access security software, infrastructure, and architectures over protected information assets
- **CC9.2** — Manages risks associated with vendors and business partners

### HIPAA-SECURITY

- **164.312(a)(1)** — Access Control
- **164.308(a)(5)(ii)(D)** — Password Management
- **164.312(d)** — Person or Entity Authentication

### ISO-27001-2022

- **A.5.21** — Managing information security in the ICT supply chain
- **A.5.23** — Information security for use of cloud services

## MITRE ATT&CK Mapping

| Technique ID | Technique Name | Tactic |
|---|---|---|
| T1078 | Valid Accounts | Defense-Evasion |
| T1098 | Account Manipulation | Persistence |
| T1078.004 | Cloud Accounts | Defense-Evasion |
| T1550.001 | Application Access Token | Defense-Evasion |
| T1528 | Steal Application Access Token | Credential-Access |

| Technique ID | Technique Name | Tactic |
|---|---|---|
| T1530 | Data from Cloud Storage | Collection |

## Sources

| Source | URL | Tier |
|---|---|---|
| **Blog** | https://www.crowdstrike.com/en-us/blog/crowdstrike-announces-contin... | T3 |
| | https://www.crowdstrike.com/en-us/blog/crowdstrike-announces-contin... | T3 |
| | https://www.pymnts.com/cybersecurity/2026/crowdstrike-launches-cont... | T3 |
| | https://cybermagazine.com/news/crowdstrike-secures-ai-agents-with-r... | T3 |
| **CrowdStrike Falcon AIDR: AI Detection & Response** | https://www.crowdstrike.com/en-us/platform/falcon-aidr-ai-detection... | T3 |



Generated 2026-06-17 13:59 UTC by TJS Security Command Center