INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-06-21 18:39 UTC

AI Agent Identity Gap: CrowdStrike Introduces Continuous Per-Action Authorization for Autonomous Workloads

SECURITY ANALYSIS | MEDIUM | CVSS 5.0

SCC Item ID	SCC-STY-2026-0239
Type	Security Analysis
Severity	MEDIUM
CVSS Base Score	5.0
Affected Products	Enterprise IAM/PAM systems (generic), CrowdStrike Falcon Next-Gen Identity Security, CrowdStrike Falcon AIDR, Falcon Zero Trust Access, AWS cloud infrastructure with standing permissions
Discovery Source	Rss:T1 Threatintel

Executive Summary

CrowdStrike has announced a capability called Continuous Identity for AI Agents, backed by its acquisition of SGNL, that replaces static session-based authentication with real-time, per-action authorization for autonomous AI workloads. The announcement surfaces a structural gap in enterprise identity architecture: current IAM and PAM controls were designed for human login events, not for non-human identities that execute thousands of actions autonomously between any human review. As agentic AI deployments accelerate across cloud environments, organizations carrying standing permissions for AI workloads face an attack surface that operates at machine speed, well beyond the reach of conventional governance controls.

Technical Analysis

The core vulnerability class CrowdStrike is addressing is not a discrete software flaw but an architectural mismatch between legacy identity models and the behavioral profile of autonomous AI agents. Traditional IAM systems authenticate a principal at session initiation and then grant that session a static permission set for its duration. For human users, this model carries acceptable risk because human action velocity is limited and observable. For AI agents, the same model creates a fundamentally different exposure profile: a single authenticated session can trigger hundreds of downstream API calls, cloud resource interactions, and data operations with no intermediate authorization check.

CrowdStrike identifies several intersecting weakness classes that define this problem. CWE-269 (Improper Privilege Management) and CWE-732 (Incorrect Permission Assignment) capture the structural issue of over-permissioned service accounts and agent identities that accumulate standing access beyond operational necessity. CWE-284 (Improper Access Control) and CWE-306 (Missing Authentication for Critical Function) describe the enforcement gap at the action level, where individual operations performed by an authenticated agent are never independently validated. CWE-287 (Improper Authentication) covers the session-initiation model's failure to persist identity assurance across an agent's operational lifetime.

The MITRE ATT&CK techniques associated with this exposure pattern are instructive. T1078 (Valid Accounts) and T1078.004 (Valid Accounts: Cloud Accounts) represent the most direct threat vector: an adversary who compromises an AI agent's identity inherits its full standing permission set. T1134 (Access Token Manipulation) and T1134.001 (Token Impersonation/Theft) describe how that compromised identity can be extended or forged. T1528 (Steal Application Access Token) and T1550.001 (Use Alternate Authentication Material: Application Access Token) cover the token-level abuse that replaces credential theft in cloud-native environments. T1548 (Abuse Elevation Control Mechanism), T1530 (Data from Cloud Storage), and T1098 (Account Manipulation) round out the post-compromise action set available to an adversary operating inside an over-permissioned agent context. T1606 (Forge Web Credentials) captures the risk at the credential-generation layer.

CrowdStrike's response, built on the SGNL acquisition, introduces a just-in-time, per-action authorization model that evaluates each agent action against current policy at execution time rather than relying on permissions granted at session start. This is architecturally consistent with zero trust principles as articulated in NIST SP 800-207, and it extends those principles to non-human identities explicitly. The announcement integrates with CrowdStrike Falcon Next-Gen Identity Security, Falcon AIDR, and Falcon Zero Trust Access, targeting AWS cloud environments with standing permissions as a primary remediation target.

The broader industry implication is significant. Most enterprise identity programs have not yet developed NHI governance frameworks that account for agent action velocity. Organizations deploying agentic AI in production, particularly those using AWS IAM roles with long-lived credentials or service accounts with broad cloud permissions, are operating with an unaudited attack surface. CrowdStrike's acquisition of SGNL signals that fine-grained, context-aware access control for non-human principals is moving from research concept to production product, and that the identity security market is reorganizing around this gap.

Action Checklist

1. Step 1: Assess NHI exposure, inventory all non-human identities in your environment, including AI agents, service accounts, automation pipelines, and API keys; identify which carry standing cloud permissions in AWS or other cloud providers (supports CIS 1.1: Establish and Maintain Detailed Enterprise Asset Inventory; CIS 5.1: Establish and Maintain an Inventory of Accounts)
2. Step 2: Audit permission scope, for each NHI identified, compare granted permissions against documented operational requirements; flag any identity with permissions exceeding the minimum necessary for its defined function (supports NIST AC-6: Least Privilege; NIST AC-3: Access Enforcement; CIS 5.4: Restrict Administrator Privileges to Dedicated Administrator Accounts)
3. Step 3: Eliminate standing permissions where possible, convert long-lived cloud credentials to short-lived, just-in-time tokens; for AWS environments, review IAM role trust policies and session duration settings; prioritize agent identities with access to data stores or administrative APIs (supports NIST AC-6: Least Privilege; D3-CRO: Credential Rotation; D3-CH: Credential Hardening)

4. Step 4: Implement behavioral monitoring for agent identities, configure audit logging to capture API calls, token issuance events, and permission escalation attempts attributed to NHI principals; establish baselines and alert on deviations (supports NIST AU-2: Event Logging; NIST AU-6: Audit Record Review, Analysis, and Reporting; CIS 8.2: Collect Audit Logs; D3-LAM: Local Account Monitoring)
5. Step 5: Update threat model with NHI-specific attack paths, map T1078.004, T1528, T1134, and T1550.001 to your cloud environment's agent identities; document the blast radius for each over-permissioned agent if compromised
6. Step 6: Evaluate CrowdStrike and SGNL capability fit, assess whether Falcon Next-Gen Identity Security or equivalent per-action authorization tooling addresses gaps identified in your NHI inventory; include this in your next IAM/PAM program review cycle
7. Step 7: Brief security leadership, present NHI inventory findings with specific risk context: number of over-permissioned agent identities, associated cloud resources accessible, and estimated time to detect compromise under current monitoring coverage

IR / Forensic Enrichment

Triage Priority	STANDARD
Escalation Criteria	Escalate to urgent if the NHI inventory reveals any agent identity with active standing credentials that have been used outside expected operating hours, from unexpected source IPs, or against resource types inconsistent with the agent's documented function — any of these patterns may indicate the structural IAM gap described in this advisory is already being exploited in your environment.
Recovery Notes	After revoking standing credentials and transitioning agent workloads to short-lived STS tokens, verify recovery by confirming no long-lived access keys remain active for NHI principals via `aws iam generate-credential-report` and validating that CloudTrail shows only `AssumeRole`-sourced authentication for agent principals going forward. Monitor CloudTrail for `iam:CreateAccessKey` or `iam:AttachUserPolicy` events attributed to agent role ARNs for a minimum of 30 days post-remediation, as adversaries who obtained standing credentials before rotation may attempt re-establishment. Rerun the `aws iam simulate-principal-policy` blast-radius assessment 30 days post-remediation to confirm permission scope reduction held through any infrastructure changes or new agent deployments.

Forensic Artifacts	AWS CloudTrail management event logs filtered on `userIdentity.type = AssumedRole` for all NHI principal ARNs — specifically `sts:AssumeRole`, `iam:CreateAccessKey`, `iam:AttachRolePolicy`, and `iam:PassRole` events, which indicate standing-credential abuse or privilege escalation by an agent identity | AWS IAM credential report (`aws iam generate-credential-report` CSV output) capturing `access_key_1_last_used_date`, `access_key_2_last_used_date`, and `password_last_used` for all service accounts and NHI principals — evidence of long-lived key usage patterns inconsistent with JIT token architecture | AWS CloudTrail data events for S3 (`GetObject`, `PutObject`, `DeleteObject`) and Lambda (`Invoke`) attributed to agent principal ARNs — reveals data access scope during a potential compromise window and supports blast-radius scoping | AWS Config configuration history for IAM roles associated with agent identities — documents when trust policies or attached managed policies were modified, identifying unauthorized permission expansion that could indicate an agent operating outside its defined authorization boundary | STS session token issuance records from CloudTrail (`GetSessionToken`, `AssumeRoleWithWebIdentity`) cross-referenced against the agent's expected invocation schedule and source context — anomalous issuance times, unexpected `aws:SourceIp` values, or unusually long `MaxSessionDuration` values are primary indicators of NHI credential misuse under the agentic AI threat model described in this advisory

Per-Action IR Details

Step 1: Assess NHI exposure — inventory all non-human identities in your environment, including AI agents, service accounts, automation pipelines, and API keys; identify which carry standing cloud permissions in AWS or other cloud providers (supports 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: Establish IR Capability and Asset Visibility

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-users`, `aws iam list-roles`, and `aws iam list-service-specific-credentials --user-name ` across all accounts to enumerate NHIs. Pipe output through `jq` to extract identities with attached policies. For on-prem service accounts, query Active Directory with `Get-ADServiceAccount -Filter * | Select-Object Name, SamAccountName, LastLogonDate` via PowerShell. Maintain results in a spreadsheet versioned in git.

Evidence: This is a preparation-phase inventory step that does not alter live state; no volatile capture is required before execution. Document the output as a baseline snapshot — AWS CloudTrail `ListRoles` and `ListUsers` API call history and any IAM Access Analyzer findings serve as the pre-existing evidentiary baseline against which future anomalous NHI activity will be measured.

Step 2: Audit permission scope — for each NHI identified, compare granted permissions against documented operational requirements; flag any identity with permissions exceeding the minimum necessary for its defined function (supports NIST AC-6: Least Privilege; NIST AC-3: Access Enforcement; CIS 5.4: Restrict Administrator Privileges to Dedicated Administrator Accounts)

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Identify and Reduce Attack Surface Prior to Incident

Controls: NIST AC-6 (Least Privilege), NIST AC-3 (Access Enforcement), CIS 5.4 (Restrict Administrator Privileges to Dedicated Administrator Accounts)

Compensating: Use AWS IAM Access Analyzer's policy validation and the AWS-native `aws iam get-account-authorization-details` command to export all inline and managed policies for NHI principals. Pipe through a Python script comparing `Action` and `Resource` fields against a documented allowlist. Flag any agent role with `*` wildcards on `Action` or `Resource`, administrative API families (`iam:*`, `s3:*`, `ec2:*`), or cross-account trust

relationships not in the operational baseline.

Evidence: Preparation-phase audit step; does not alter live state. Preserve the raw IAM policy JSON output and Access Analyzer findings as a dated artifact before any permission changes are made — this establishes the pre-remediation permission posture needed to demonstrate scope of over-permissioning if a compromise is later discovered.

Step 3: Eliminate standing permissions where possible — convert long-lived cloud credentials to short-lived, just-in-time tokens; for AWS environments, review IAM role trust policies and session duration settings; prioritize agent identities with access to data stores or administrative APIs (supports NIST AC-6: Least Privilege; D3-CRO: Credential Rotation; D3-CH: Credential Hardening)

NIST Phase: Containment

Reference: NIST 800-61r3 §3.3 — Containment Strategy: Reduce Blast Radius of Compromised NHI Credentials

Controls: NIST AC-6 (Least Privilege), NIST AC-12 (Session Termination)

Compensating: Before revoking any standing credentials, capture a snapshot of active AWS STS sessions via `aws sts get-caller-identity` and active session tokens via `aws iam list-service-specific-credentials`. Then use `aws iam delete-access-key` for long-lived keys and reconfigure agent workloads to assume IAM roles via `sts:AssumeRole` with `MaxSessionDuration` set to 900–3600 seconds. For agents accessing S3 data stores or admin APIs, enforce resource-level conditions in the trust policy restricting `aws:SourceIp` or `aws:PrincipalTag`.

Evidence: BEFORE revoking standing credentials or modifying trust policies, capture: (1) AWS CloudTrail `GetSessionToken`, `AssumeRole`, and `AssumeRoleWithWebIdentity` events for each NHI principal over the past 90 days to establish usage baseline and detect any anomalous cross-account or cross-region calls; (2) current STS active session list; (3) any active API Gateway or Lambda invocation logs showing the agent's request patterns. Revoking credentials destroys session context and may sever forensic continuity if an agent is already compromised.

Step 4: Implement behavioral monitoring for agent identities — configure audit logging to capture API calls, token issuance events, and permission escalation attempts attributed to NHI principals; establish baselines and alert on deviations (supports NIST AU-2: Event Logging; NIST AU-6: Audit Record Review, Analysis, and Reporting; NIST SI-4 referenced as a monitoring control — note: SI-4 is not in the provided knowledge base extract; cite only AU-2 and AU-6 from verified data; CIS 8.2: Collect Audit Logs; D3-LAM: Local Account Monitoring)

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis: Monitor NHI Principals for Behavioral Anomalies

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

Compensating: Enable AWS CloudTrail in all regions with data events for S3 and Lambda, then export logs to an S3 bucket with CloudWatch Logs integration. Write CloudWatch Metric Filters targeting `userIdentity.type = AssumedRole` AND `userIdentity.principalId` matching your NHI role ARN pattern, alerting on `sts:AssumeRole` chains exceeding two hops, `iam:CreateAccessKey` or `iam:AttachRolePolicy` called by an agent principal, and API call volume spikes >3 standard deviations from the 30-day baseline. Sigma rule `aws_iam_privilege_escalation_via_policy` (SigmaHQ community ruleset) can be adapted for this purpose.

Evidence: Capture before altering any logging configuration: existing CloudTrail trail configuration (`aws cloudtrail describe-trails`), current CloudWatch log group retention settings, and a point-in-time export of recent NHI API activity. Key event sources to baseline: CloudTrail management events for `sts:AssumeRole`, `iam:PassRole`, `iam:CreateAccessKey`; CloudTrail data events for S3 `GetObject`/`PutObject` by agent principals; and AWS Config rule evaluation history for IAM policy changes attributed to NHI accounts.

Step 5: Update threat model with NHI-specific attack paths — map T1078.004, T1528, T1134, and T1550.001 to your cloud environment's agent identities; document the blast radius for each over-permissioned agent if compromised

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Threat Modeling and Attack Path Documentation for NHI Principals

Controls: NIST AC-6 (Least Privilege), NIST AC-4 (Information Flow Enforcement)

Compensating: For each over-permissioned agent role, run `aws iam simulate-principal-policy --policy-source-arn --action-names iam:CreateAccessKey sts:AssumeRole s3:GetObject` to enumerate reachable actions. Build a blast-radius matrix in a spreadsheet: columns are accessible resource types (S3 buckets, RDS, EC2, cross-account roles), rows are agent identities, cells are data classification and estimated exfiltration volume. This replaces expensive commercial attack-path tooling for small teams.

Evidence: This step does not alter live state and requires no volatile capture. Preserve the `simulate-principal-policy` output and IAM policy JSON as dated artifacts to support future tabletop exercises or post-incident reviews. If a prior incident or anomaly triggered this threat model update, retain the CloudTrail events that surfaced the concern as evidentiary context for the modeling assumptions.

Step 6: Evaluate CrowdStrike and SGNL capability fit — assess whether Falcon Next-Gen Identity Security or equivalent per-action authorization tooling addresses gaps identified in your NHI inventory; include this in your next IAM/PAM program review cycle

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: Program Improvement and Capability Gap Closure

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

Compensating: Without budget for Falcon or SGNL, implement NIST-aligned compensating controls: (1) AWS SCPs (Service Control Policies) at the Organizations level to enforce hard deny on `iam:CreateAccessKey` by non-human principals outside approved automation accounts; (2) open-source `parliament` library (Duo Security) for continuous IAM policy linting in CI/CD; (3) scheduled Lambda function invoking `aws iam generate-credential-report` weekly and publishing findings to a security team SNS topic.

Evidence: No live-state alteration in this step; no volatile capture required. Document the gap analysis output — the delta between capabilities identified in the NHI inventory (Steps 1–2) and what current tooling can monitor or enforce — as the formal input artifact for the IAM/PAM program review. This artifact supports audit evidence requirements under NIST AU-11 (Audit Record Retention) if the review is part of a compliance-driven assessment cycle.

Step 7: Brief security leadership — present NHI inventory findings with specific risk context: number of over-permissioned agent identities, associated cloud resources accessible, and estimated time to detect compromise under current monitoring coverage

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: Lessons Learned Reporting and Leadership Communication

Controls: NIST AU-6 (Audit Record Review, Analysis, and Reporting), CIS 7.2 (Establish and Maintain a Remediation Process)

Compensating: Construct the briefing from outputs already produced: NHI inventory spreadsheet (Step 1), permission audit flags (Step 2), blast-radius matrix (Step 5), and CloudTrail baseline gap analysis (Step 4). Quantify mean-time-to-detect by replaying a simulated `sts:AssumeRole` chain in a non-production account and measuring how long before a CloudWatch alert fires — or confirm no alert fires, which is itself the key risk metric for leadership.

Evidence: No live-state alteration; no volatile capture required. The briefing package itself — NHI inventory, permission audit findings, and detection gap quantification — constitutes a structured post-incident artifact. Retain it under your IR documentation retention policy (NIST AU-11) to support future audits or regulatory inquiries, particularly if any of the inventoried agent identities subsequently appear in a security incident.

Detection Guidance

Detection for this threat class focuses on anomalous behavior from non-human identity principals rather than static indicators of compromise. Key detection surfaces include:

Cloud IAM and API logs: Monitor AWS CloudTrail (or equivalent) for NHI principals generating unusual API call volumes, accessing resources outside their established operational baseline, or requesting permissions beyond their assigned role scope. Alert on any token issuance or role assumption (sts:AssumeRole) event where the requesting principal is an AI agent or automation service account (relevant to T1078.004, T1528, T1550.001).

Token manipulation signals: Hunt for patterns consistent with T1134 and T1134.001: access tokens used from unexpected source IPs or regions, tokens reused after documented session expiration, or token issuance events not preceded by expected authentication flows.

Privilege escalation attempts: Log and alert on any NHI principal attempting to modify its own IAM policies, attach new policies, or create new credentials (T1548, T1098). These actions fall outside the operational scope of well-governed agent identities.

Data access anomalies: Flag NHI principals accessing cloud storage buckets (T1530) or databases not referenced in their documented function. Volume-based anomalies, such as bulk reads or exports, warrant immediate investigation.

Authentication gaps: Audit for AI agent sessions where no per-action authorization checkpoint exists between initial authentication and downstream operations. This is a policy audit rather than a log-based detection, and it surfaces the structural CWE-306 exposure directly.

Framework alignment: NIST AU-2 (Event Logging) and AU-6 (Audit Record Review, Analysis, and Reporting) provide the logging foundation. CIS 8.2 (Collect Audit Logs) establishes the baseline collection requirement. D3-LAM (Local Account Monitoring) and D3-UAP (User Account Permissions) map to the behavioral monitoring and permission review activities described above.

Framework Mappings

MITRE-ATTACK

- **T1606** — Forge Web Credentials
- **T1134** — Access Token Manipulation
- **T1134.001** — Token Impersonation/Theft
- **T1528** — Steal Application Access Token
- **T1548** — Abuse Elevation Control Mechanism
- **T1530** — Data from Cloud Storage
- **T1550** — Use Alternate Authentication Material
- **T1098** — Account Manipulation
- **T1550.001** — Application Access Token
- **T1078** — Valid Accounts
- **T1078.004** — Cloud Accounts

NIST-800-53R5

- **AC-6** — Least Privilege
- **CM-6** — Configuration Settings
- **AC-2** — Account Management
- **IA-2** — Identification and Authentication (Organizational Users)

- **IA-5** — Authenticator Management
- **AC-3** — Access Enforcement
- **IA-8** — Identification and Authentication (Non-Organizational Users)

OWASP-TOP10-2021

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

CIS-V8

- **3.3** — Configure Data Access Control Lists
- **6.1** — Establish an Access Granting Process
- **6.2** — Establish an Access Revoking Process
- **6.3** — Require MFA for Externally-Exposed Applications
- **6.4** — Require MFA for Remote Network Access
- **6.5** — Require MFA for Administrative Access
- **5.4** — Restrict Administrator Privileges to Dedicated Administrator Accounts
- **6.8** — Define and Maintain Role-Based Access Control

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
- **164.312(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
T1606	Forge Web Credentials	Credential-Access
T1134	Access Token Manipulation	Defense-Evasion
T1134.001	Token Impersonation/Theft	Defense-Evasion
T1528	Steal Application Access Token	Credential-Access
T1548	Abuse Elevation Control Mechanism	Privilege-Escalation

Technique ID	Technique Name	Tactic
T1530	Data from Cloud Storage	Collection
T1550	Use Alternate Authentication Material	Defense-Evasion
T1098	Account Manipulation	Persistence
T1550.001	Application Access Token	Defense-Evasion
T1078	Valid Accounts	Defense-Evasion
T1078.004	Cloud Accounts	Defense-Evasion

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://cybermagazine.com/news/ai-cyber-attacks-risk-tech-this-week...	T3
	https://www.pymnts.com/cybersecurity/2026/crowdstrike-launches-cont...	T3
CrowdStrike Falcon AIDR: AI Detection & Response	https://www.crowdstrike.com/en-us/platform/falcon-aidr-ai-detection...	T3

Generated 2026-06-21 18:39 UTC by TJS Security Command Center