AI & Agentic Risk Posture Assessment

ASSESSMENT FINDINGS

What Sonrai found across your cloud estate

219 accounts · AWS · Full cloud organization · 8 threat vectors · Read-only

> **Mythos finds the vulnerability. Your permissions posture determines the damage.**
> **Sonrai research · Apr 2026** — industry researchers predict ~6 months until open-weight models reach Mythos-class vulnerability discovery.

MYTHOS-READY ASSESSMENT

ASSESSMENT FINDING

AI agents are operating in your cloud with far more permission than they need.

28 AI identities operate in your estate — **every one holds privileged access** across 35 AI/ML services. Sonrai's engine flags **6 of 8 threat vectors as AI-applicable** — Compliance Bypass, Ransomware, Network Bypass, Logging Evasion, DNS Manipulation, Execution Control. Standing permissions AI agents inherit remain invisible and uncontrolled at the access plane.

Assessment Results

IAM & SERVICE CONTROL PLANE

219 accounts · AWS · Full org scan

AI IDENTITIES DETECTED

Across **SSO Users, Agents and Workload Identities**

28
AI Identities Discovered

• UNMANAGED SURFACE
Spans **SSO Users, Agents, and Workload Identities** — grows weekly with Claude Code & Copilot adoption.

28
Over-Privileged

• BLAST RADIUS
100% of your AI identities hold privileged access — every one can write or delete production data beyond their task. One prompt injection or leaked key = **instant breach**.

35
AI/ML Services Deployed

• ATTACK SURFACE
Bedrock, SageMaker, Textract, Amazon Q, Rekognition, Forecast, Personalize, and 28 others — every deployed AI/ML service is a fresh attack path an agent can invoke.

12
Admin-Equivalent

• CLOUD TAKEOVER
12 AI identities hold near-admin access — chaining **sts:AssumeRole → iam:CreateRole** or equivalent. Every one is a potential cloud-takeover path.

5,299	1,754	179	3,400
Excessive Privileges	**Zombie Identities**	**Unused Services Exposed**	**Evasion-Capable Identities**
Standing access · never used	Dormant 90+ days · full permissions	Standing access · never invoked	Bypass logging · compliance · DNS

THREAT VECTORS — PERMISSIONS & IDENTITIES AT RISK ACROSS YOUR ESTATE

6 of 8 **Threat Vectors Carry AI Risk**
Flagged by Sonrai's own engine — identities on these paths are directly reachable by AI agents

SONRAI AI CLASSIFIER

Compute Hijacking — CRITICAL
3,600 identities · 30 permissions
Cloud cost abuse, performance degradation, exploitation of trusted IP space.

Compliance & Audit Bypass — AI · CRITICAL
2,800 identities · 68 permissions
Missed violations, failed audits, regulatory fines.

Data Destruction & Ransomware — AI · CRITICAL
2,300 identities · 25 permissions
Operational downtime, financial extortion, irreversible data loss.

Identity Takeover & Escalation — CRITICAL
2,100 identities · 109 permissions
Data theft, system disruption, breach of customer trust.

Network Security Bypass — AI · HIGH
1,800 identities · 69 permissions
Exposure of internal systems, bypass of segmentation controls.

Logging & Monitoring Evasion — AI · HIGH
1,600 identities · 63 permissions
Longer dwell time, delayed response, undermined forensics.

DNS & Traffic Manipulation — AI · HIGH
1,300 identities · 40 permissions
Credential theft, data leakage, service impersonation.

Execution & Command Control — AI · HIGH
1,200 identities · 41 permissions
Automated lateral movement, exfiltration, infrastructure sabotage.

THIRD-PARTY EXPOSURE

38 **External integrations** with standing IAM access

0 currently restricted — every vendor integration requires owner review & verification.

DORMANT IDENTITIES

1,754 **Zombie identities** unused 90+ days

2 are AI-labeled zombies. The rest are the service-account credentials AI agents inherit when invoked.

AI/ML ATTACK SURFACE

35 **AI/ML services deployed** reachable by your 28 AI identities

Bedrock, SageMaker, Textract, Amazon Q, Rekognition + 30 others.

WHEN MYTHOS LANDS IN YOUR ESTATE

T – 6 months · Industry prediction

TODAY (PRE-MYTHOS)
An over-privileged identity is a **latent** risk — exploited only if an attacker finds a vulnerability they can reach. Most stay undiscovered for months.

→

WITH MYTHOS-CLASS AI
Discovery becomes **automated, continuous**. Your **28 privileged AI identities** × **35 AI/ML services** become always-on blast radius — no longer latent.

REC SEE IT LIVE

Watch a developer's Claude Code session accidentally go sideways — and Sonrai's guardrails catch it in real time.

Recorded walkthrough — real AWS account, real agent, real-time enforcement at the permission layer.

SCAN TO WATCH

YOUR REMEDIATION PLAN

Mythos is ~6 months out — shrink the blast radius now.

3 sequenced steps · Same-day execution · **Remaining work**

EVERY AGENTIC ACCESS VECTOR · PROTECTED

AI agents provision infrastructure without human review. CPF enforces permissions at the access plane — wherever they reach cloud.

VS Code · Codex · Claude · Cursor · SSH · Console · Web / SaaS · AI Agents · Terraform · CI/CD

COMPLIANCE COVERAGE: SOC 2 · ISO 27001 · PCI-DSS 4.0 · NIST 800-53

YOUR POSTURE · BEFORE → AFTER

One action cycle · Same-day execution · Fully reversible

~~5,299~~ → 0	~~179~~ → 0	~~357~~ → 357	~~1,754~~ → 0	~~3,400~~ → 0
Over-Privileged Identities	**Unused Services** Exposed	**Active Services** Unprotected → Protected	**Zombie** Identities	**Evasion-Capable** Identities

179 — Services Disabled

STEP 1 · SAME DAY

- Disable **179 unused services** across all 219 accounts — removing standing access AI agents inherit
- Closes the Compliance Bypass, Ransomware, and Network Bypass vectors at the service layer

1,754 — Zombies Quarantined

STEP 2 · SAME DAY

- Quarantine **1,754 dormant identities** — unused 90+ days, credentials AI agents most often inherit
- Lock **38 third-party integrations** to owner-verified only — currently 0 are restricted
- Any active identity gets reinstated via JIT in seconds

5,299 — Identities Right-Sized

STEP 3 · ONGOING

- Strip unused permissions from **5,299 over-privileged identities** — the standing access driving every threat vector
- Right-size the **357 protected services** — restrict to operations actually used
- **Default-deny every new AI agent** — zero standing permission from day one
- Neutralizes all **8 threat vectors** mapped above — 6 of which carry direct AI risk

AGENTIC GUARDRAILS · ONGOING PROTECTION

For the services your business **does** want active

357

Services Actively Protected

Every privileged permission on your remaining active services — wrapped in **default-deny + just-in-time authorization**. AI agents can't invoke anything beyond the task.

Default-Deny for AI Agents
Every new agent or workload starts with **zero standing permission**

Just-In-Time Grant-Back
Seconds to authorize a specific action — **auto-revoked after use**

Per-Permission Inspection
Block **ransomware, exfil, escalation** at the permission layer

JIT APPROVALS ROUTE VIA Slack · Microsoft Teams

PROJECTED VALUE

Zero — **Standing access for AI agents**
Day 1 · zero trust on every new identity

- **Default-deny every new AI agent** at the moment it's detected — no permission sprawl, no standing privilege inherited from a human identity.
- **6 of 8 threat vectors** carrying direct AI risk — neutralized at the access plane in real time.
- **7,232 risk items** eliminated in one action cycle — 5,299 over-privileged identities, 1,754 zombies, 179 unused services.
- **Audit evidence auto-generated** for SOC 2 CC6.1, NIST AC-6, PCI-DSS 7 — continuous compliance, not quarterly scramble.

WHAT CPF CATCHES IN PRACTICE

Compromised AI identity chains `sts:AssumeRole` into a backup account
Blocked at the access plane in <50ms · **SOC alerted**

Claude Code agent requests write to a customer-data S3 bucket
JIT routed to approver in Slack · denied in **8 seconds** · zero standing grant

A 90-day zombie role is invoked unexpectedly
Auto-quarantine triggers · identity locked pending owner review