SHADOW AI
DISCOVERY · SECSTUDIO

SHADOW AI — FLEET DISCOVERY REPORT

AI Bill of Materials & Fleet Discovery

Prepared for **Northwind Trading Co.**

AS OF	FLEET SCOPE	CLASSIFICATION
12 Aug 2026 · continuous	**320 endpoints monitored**	**Confidential**
AI ARTIFACTS	REPORT ID	AI RISK POSTURE
1,284 discovered	**SAD-FLEET-2026-0142**	**64 / 100**

SECTION 1

Fleet overview

41%

ENDPOINTS RUNNING UNSANCTIONED AI

131 of 320 endpoints

320

Endpoints monitored

314 / 320

Sensors online

1,284

AI artifacts discovered

268

Endpoints with AI present

The endpoint sensor is deployed across **320 endpoints**, of which **314 are reporting** (6 offline > 180s). It has inventoried **1,284 AI artifacts** on **268 endpoints** — 84% of the fleet has some form of AI installed. Of those, **41% run unsanctioned AI** outside approved policy: local LLM runtimes, coding agents, browser extensions and unregistered MCP servers. Six sensors require attention (2 offline, 4 stale heartbeat).

OPERATING-SYSTEM BREAKDOWN & SENSOR HEALTH

PLATFORM	ENDPOINTS	SENSORS ONLINE	WITH AI	ARTIFACTS	HEALTH
Windows 10 / 11	190	187 / 190	158	742	Degraded
macOS 13–15	96	95 / 96	84	398	Healthy
Linux (Ubuntu / RHEL)	34	32 / 34	26	144	Degraded
Fleet total	**320**	**314 / 320**	**268**	**1,284**	98% Online

DISCOVERY CADENCE & FRESHNESS

Network scan	every 15 s
Static inventory sweep	every 600 s
Sensor heartbeat	every 60 s
Marked offline after	180 s no heartbeat

Detections combine a **real-time network collector** (process, socket and egress signals every 15s) with a slower **static sweep** (installed packages, extensions, config and key material every 10min). A missed heartbeat past **180s** flags the endpoint offline and excludes its inventory from freshness-weighted counts until it re-checks in. Current mean inventory age: **4m 12s**.

SECTION 2

AI artifacts by category

All **1,284 discovered artifacts** classified by detection category. Distribution below is weighted by artifact count; badges reflect the category's dominant risk contribution to fleet posture.

DISCOVERY DISTRIBUTION

LLM apps / SDKs **512** · Coding agents **173** · API keys / secrets **96** · MCP servers **88** · Browser AI extensions **214** · CLI AI tools **201**

CATEGORY DETAIL

DETECTION CATEGORY	COUNT	ENDPOINTS	EXAMPLE ARTIFACTS	RISK
LLM apps / SDKs	512	231	ChatGPT desktop, Claude desktop, Ollama local models (llama3, mistral), `openai` / `anthropic` SDKs	High
Coding agents	173	61	Unapproved AI coding assistants & IDE agents with repo + credential access	High
MCP servers	88	44	37 registered / 51 unregistered · 12 on vulnerable versions incl. SHAI-2026-001	Critical
API keys / secrets	96	38	OpenAI & Anthropic keys in LLM prompts, `.env` files, shell history and MCP configs	Critical
Browser AI extensions	214	119	Sidebar AI assistants, autofill/writing helpers with page-content egress	Medium
CLI AI tools	201	88	Terminal LLM clients, AI git helpers, prompt runners & local inference CLIs	Medium
Total	**1,284**	**268***	*distinct endpoints; many carry artifacts in multiple categories	

The two smallest categories by volume — **MCP servers (88)** and **API keys / secrets (96)** — carry the highest risk weight. 51 of 88 MCP servers are unregistered and 12 run known-vulnerable versions; 96 live secrets were recovered from prompts and configs, the primary driver of the fleet's critical findings.

SECTION 3

AI Bill of Materials (AIBOM)

Component-level inventory of discovered AI software. **Provenance** marks whether the publisher and package signature were verified against a trusted source. Risk badge reflects the component's fleet exposure.

ARTIFACT	TYPE	VERSION	ENDPOINTS	SOURCE / PUBLISHER	PROVENANCE	RISK
ChatGPT Desktop	LLM app	`1.2026.194`	88	OpenAI	Verified	Medium
Claude Desktop	LLM app	`0.9.3`	54	Anthropic	Verified	Low
Ollama + local models	Local runtime	`0.3.9`	37	Ollama	Verified	High
llama3-8b (egress observed)	Local model	`q4_0`	11	self-hosted	Unverified	High
openai (Python SDK)	LLM SDK	`1.3.5` (EOL)	46	PyPI	Unverified	High
anthropic (Node SDK)	LLM SDK	`0.27.0`	29	npm	Verified	Low
server-filesysem (typosquat)	MCP server	`0.1.2`	3	npm (untrusted)	Unverified	Critical
Unregistered MCP server SHAI-2026-001	MCP server	`0.4.1` (vuln)	6	unknown	Unverified	Critical
mcp-server-git	MCP server	`2026.6.0`	14	modelcontextprotocol	Verified	Medium
AI coding assistant (unapproved)	Coding agent	`2.14.0`	31	3rd-party marketplace	Unverified	High
Browser AI sidebar ext.	Browser ext.	`4.7.1`	48	Chrome Web Store	Unverified	Medium
llm (CLI client)	CLI tool	`0.15`	27	PyPI	Verified	Low
ai-commit git helper	CLI tool	`1.9.0`	22	npm	Unverified	Medium
whisper.cpp local STT	Local model	`1.6.2`	9	GitHub	Unverified	Medium

Highlighted rows are the fleet's supply-chain concerns: the typosquatted **server-filesysem** MCP package, the unregistered vulnerable **SHAI-2026-001** server, an **end-of-life openai SDK** on 46 endpoints, and a **local llama3 model with observed network egress** on 11 endpoints. All four are unverified provenance and are staged for containment in Section 5.

SECTION 4

Top risky endpoints & identities

Highest-risk endpoints ranked by composite of artifact count, provenance and observed data-egress signals.

HOSTNAME	OS	ARTIFACTS	TOP RISK	SEVERITY
NWT-DEV-014	Windows 11	23	OpenAI key in ChatGPT prompt + unregistered MCP server	Critical
NWT-ENG-007	macOS 15	19	server-filesysem typosquat + Ollama egress	Critical
NWT-DEV-022	Linux (Ubuntu)	18	SHAI-2026-001 vulnerable MCP server	Critical
NWT-FIN-003	Windows 10	14	Anthropic key in .env + shadow browser AI ext.	High
NWT-DEV-031	macOS 14	13	Unapproved coding agent w/ repo + cred access	High
NWT-ENG-019	Windows 11	12	EOL openai SDK 1.3.5 + local llama3 egress	High
NWT-DAT-005	Linux (RHEL)	11	Key in shell history + 3 CLI AI tools	High
NWT-MKT-011	Windows 10	9	4 browser AI extensions w/ page egress	Medium
NWT-DEV-045	macOS 15	8	ai-commit unverified + whisper.cpp local model	Medium

IDENTITIES & OFFBOARDING

9

Offboarded users still holding active AI-tool access

Directory shows these accounts as terminated, yet the sensor still observes valid AI-tool sessions or provisioned API keys tied to them. No automated de-provisioning on offboarding.

`j.okafor`	ChatGPT Team seat active	High
`r.silva`	OpenAI API key still live	Critical
`m.kaur`	Coding-agent token valid	High
`+ 6 more`	Browser-ext / CLI sessions	Medium

Nine departed users retain a path to corporate data through AI tooling. Two hold live API keys, which are also counted among the 96 exposed secrets. Recommend immediate revocation and directory-driven de-provisioning (Section 5).

SECTION 5

Coverage, methodology & next scan

98%
FLEET SENSOR COVERAGE
314 of 320 endpoints reporting

9
Collectors per sensor

4m 12s
Mean inventory age

2
Sensors offline

4
Stale heartbeats

WHAT THE SENSOR COLLECTS (9 COLLECTORS)

1 · Process & runtime	LLM apps, local models
2 · Installed packages	SDKs, CLI tools
3 · Browser extensions	AI sidebars, helpers
4 · MCP registry & configs	registered vs not
5 · Network egress	AI endpoints, sockets
6 · Secret scan	keys in env / prompts
7 · Coding-agent detect	IDE agents, tokens
8 · Identity binding	tool ↔ directory user
9 · Provenance check	publisher + signature

PRIVACY

The sensor collects **metadata, not content**. It records the presence, version, publisher and network behavior of AI artifacts and detects the *presence* of secret patterns; it does not exfiltrate prompt text, file contents or the secret values themselves. Key material is fingerprinted (hashed prefix) for de-duplication and never stored in clear. Collection is scoped to corporate-managed endpoints.

RECOMMENDED REMEDIATION

1 **Rotate exposed secrets & enable AI-DLP** — 0–2 weeks
Rotate the 96 detected keys (2 tied to offboarded users) and block secret egress to public LLMs inline at the endpoint.

2 **Quarantine untrusted AI components** — 0–3 weeks
Isolate SHAI-2026-001, remove the server-filesysem typosquat, and upgrade the EOL openai SDK across 46 endpoints.

3 **Sync identity de-provisioning** — 2–6 weeks
Wire directory offboarding to revoke AI-tool seats, keys and agent tokens automatically for the 9 departed users.

NEXT SCHEDULED SWEEP

Continuous discovery is always on. The next full static inventory reconciliation completes within **600s**; a delta AIBOM is regenerated on any new artifact detection. Full evidence and per-endpoint records are available in the platform console and the accompanying Detailed Findings Report.