DEFENDAI™

Securing Agentic & Generative AI

AI Agent Security Report

Prepared for	**Acme**
Assessment Date	June 4, 2026
Scope	3 machines · 271 agents · Multi-machine
Tier	AgentDiscover Starter
Prepared by	DefendAI AgentDiscover Platform

CONFIDENTIAL

Executive Summary

DefendAI AgentDiscover performed a comprehensive scan of the Acme environment on June 4, 2026. The assessment covered 3 machines, identified 271 AI agents, and detected 0 credential exposures. The majority of agents are classified as unknown. 271 agents are ungoverned and require review.

271	3	0	271	0
Total Agents	Machines Scanned	Credential Exposure	Ungoverned	Governed (Approved)

Agent Classification Summary

Classification	Count	Description	Risk Level
CONFIRMED	8	Active in production with verified source code	N/A
UNKNOWN	244	Found in code but not currently running	Medium
GHOST	4	Active runtime — no source code found	Critical
SHADOW AI	15	Known apps or direct LLM API calls without governance	Medium–High
ZOMBIE	0	Previously active, now inactive	Low

Machine Inventory

Machine	Agents	Cred. Exposure	Ungoverned %	Last Scan	Status
defendai-demo	19	0	100%	Jun 4, 2026, 5:07 PM	Active
LAPTOP-IKI4CVNF	25	0	100%	May 26, 2026, 11:11 PM	Active
Omars-MacBook-Pro.local	227	0	100%	May 26, 2026, 6:16 PM	Active

Critical Findings

FINDING 1 HIGH

271 Ungoverned Agents — No Owner Assignment

Location: defendai-demo, LAPTOP-IKI4CVNF, Omars-MacBook-Pro.local

271 agents have governance_status = unreviewed. No owner has been assigned, no approval decision has been made, and no policy enforcement is in place. These agents may be making direct LLM API calls without audit logging or DLP controls.

Recommendation

Use the DefendAI governance workflow to assign owners and approve or reject each agent. Prioritize agents with network access to external LLM providers. Upgrade to AgentWatch to enable real-time prompt monitoring for approved agents.

FINDING 2 MEDIUM

CrewAI, LangChain/LangGraph, agentdiscover-scanner, AutoGen, Direct HTTP LLM Client, Shadow AI, LLM API Endpoint Detected Across 3 Machines

Location: All 3 machines

AI framework(s) detected across multiple scanned machines (CrewAI, LangChain/LangGraph, agentdiscover-scanner, AutoGen, Direct HTTP LLM Client, Shadow AI, LLM API Endpoint). Shared automation infrastructure may be routing sensitive data between systems without consistent governance.

Recommendation

Audit workflows for AI agent usage across machines. Ensure LLM API access is logged and governed. Consider routing LLM calls through DefendAI Gateway for centralized policy enforcement.

Governance Status

The following table reflects the current governance posture as of the report date. Agents are reviewed using the DefendAI governance workflow (assign owner !' approve/reject !' monitor).

Status	Count	% of Total	Action Required
Approved	0	0%	None — continue monitoring
Under Review	0	0%	Complete review and make decision
Unreviewed	271	100%	Assign owners and review immediately
No Identity Link	138	50.9%	Re-run scan to establish correlation

Remediation Priorities

Priority	Action	Timeline	Owner
P1 — Immediate	Assign owners to 248 ghost/unknown agent(s)	24–48 hours	AI Platform Team
P2 — Short Term	Review and approve/reject all 271 unreviewed agent(s)	1–2 weeks	Department Heads
P3 — Medium Term	Deploy DefendAI Gateway for LLM traffic routing	1 month	Engineering
P3 — Medium Term	Establish AI governance policy document	1 month	CISO
P3 — Medium Term	Enable continuous scanning (daemon mode)	1 month	IT Ops

Agent Inventory — Top 20 by Risk

The following agents are sorted by risk level. Full inventory available in the DefendAI platform at acme.defendai.ai/dashboard/agent_inventory.

Agent Name	Classification	Risk	Machine	Credentials	Gov. Status
shadow-agent	GHOST	Critical	defendai-demo	—	Unreviewed
ecrg_app	UNKNOWN	Critical	Omars-MacBook-Pro.local	—	Unreviewed
app_rate_limiter_mailgun_integration	UNKNOWN	Critical	Omars-MacBook-Pro.local	—	Unreviewed
evaluate	UNKNOWN	Critical	Omars-MacBook-Pro.local	—	Unreviewed
illustrator_tool	UNKNOWN	Critical	Omars-MacBook-Pro.local	—	Unreviewed
openRouterUtils	UNKNOWN	Critical	Omars-MacBook-Pro.local	—	Unreviewed
resume_agent	UNKNOWN	Critical	Omars-MacBook-Pro.local	—	Unreviewed
elchanio_wk1_lab2_llm_parallel_evaluation	UNKNOWN	Critical	Omars-MacBook-Pro.local	—	Unreviewed
scanner_agent	UNKNOWN	Critical	Omars-MacBook-Pro.local	—	Unreviewed
member	UNKNOWN	Critical	Omars-MacBook-Pro.local	—	Unreviewed
simple_client_llm	UNKNOWN	Critical	Omars-MacBook-Pro.local	—	Unreviewed
frontier_agent	UNKNOWN	Critical	Omars-MacBook-Pro.local	—	Unreviewed
agent	UNKNOWN	Critical	Omars-MacBook-Pro.local	—	Unreviewed
orchestrator_workers_demo	UNKNOWN	Critical	Omars-MacBook-Pro.local	—	Unreviewed
main	UNKNOWN	Critical	Omars-MacBook-Pro.local	—	Unreviewed
backend_api	UNKNOWN	Critical	Omars-MacBook-Pro.local	—	Unreviewed
project-planner	UNKNOWN	Critical	Omars-MacBook-Pro.local	—	Unreviewed
evaluation	UNKNOWN	Critical	Omars-MacBook-Pro.local	—	Unreviewed
traders	UNKNOWN	Critical	Omars-MacBook-Pro.local	—	Unreviewed
weather_ui	UNKNOWN	Critical	Omars-MacBook-Pro.local	—	Unreviewed

Full agent list (271 agents) available in platform. Export as AIBOM via Agent Inventory !' Export.

About This Report

This report was generated by the DefendAI AgentDiscover platform. AgentDiscover performs multi-layer AI agent discovery using static code analysis (Layer 1), passive network monitoring (Layer 2), Kubernetes eBPF monitoring via Tetragon (Layer 3), and endpoint analysis via osquery (Layer 4).

Component	Version	Detection Method
AgentDiscover Scanner	v2.4.0	Layer 1 (static), Layer 2 (network), Layer 3 (eBPF), Layer 4 (endpoint)
AgentDiscover Platform	v1.0	Correlation engine, governance workflow, multi-machine aggregation
Scan Duration	30 seconds	Network observation window per machine
Report Generated	June 14, 2026 at 2:40 AM	On-demand export from Agent Inventory

For questions about this report, contact DefendAI at info@defendai.ai or +1.408.656.9708. Platform access: acme.defendai.ai