ONSITE AI GOVERNANCE × jamf

Govern AI with confidence.

AI is being implemented faster than it can be governed. Your people are already running Claude Code, Codex and AI assistants on the Macs you manage — the question is whether you can **see it, control it, and prove it**.

81.7%

of Apple enterprises are exposed to AI risk — already lived, or actively expected.

Jamf AI Governance Survey · 687 IT & security leaders · Q2 2026

THE FINDING

AI is being implemented **faster than it can be governed.**

Nearly three-quarters of Apple enterprises have already deployed AI. The conversation has moved past whether to adopt — and risk is scaling with adoption, not against it.

72.9%
have deployed AI in some form

22.0%
have already had a cost or security incident

36.7%
name AI governance a top priority — level with enablement

Incident rate by depth of AI adoption

The counterintuitive finding: the organisations furthest along with AI report incidents at the **highest** rate.

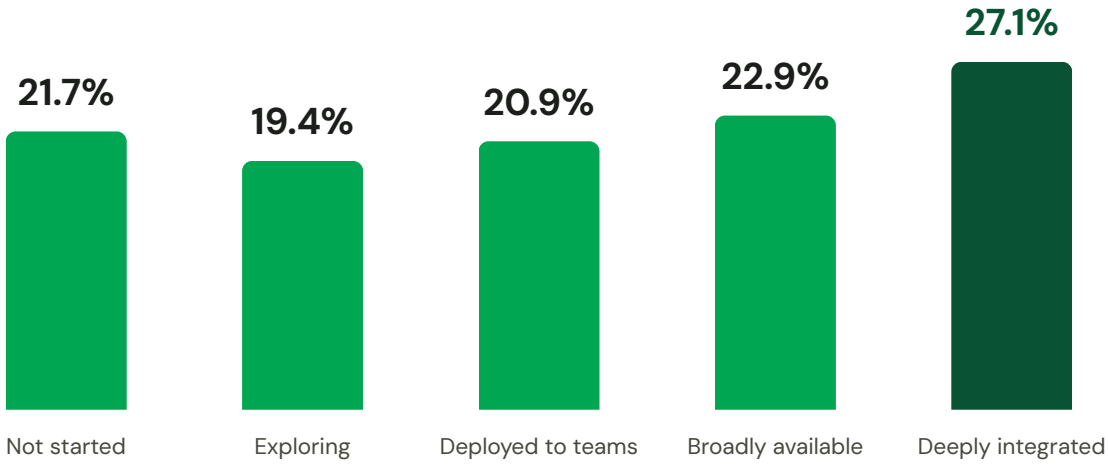

40% higher incident rate among deeply-integrated adopters (27.1%) than those still exploring (19.4%). Adoption isn't the risk — **ungoverned adoption at scale is**.

Source: Jamf AI Governance Survey, Q2 2026 (n=687).

SHADOW AI IS ALREADY INSIDE

Two ungoverned populations, **two risk surfaces.**

DEVELOPER TEAMS

Code & agents

Claude Code and Cursor bring code exfiltration, unrestricted model access, and autonomous agents acting without policy boundaries — plus invisible CLI tools, IDE extensions and packages pulled from GitHub.

BROADER WORKFORCE

Personal accounts

Staff use personal AI accounts to upload sensitive documents and query business systems through MCP connectors — with no visibility, no audit trail, and no board-level defensibility.

> Blocking known AI sites is the easy part. **CLI tools, IDE extensions and packages from GitHub are largely invisible** — when one vector gets closed, users find another.

— from 178 open-ended responses, Jamf AI Governance Survey

THE CHALLENGES CLUSTER INTO FOUR THEMES

Shadow AI

Unsanctioned tools and personal accounts; IT left in the dark about what's used.

Agentic & dev AI

Command-line tools, IDE extensions and agents with real access to code and systems.

Vendor sprawl

Every product embeds AI; vetting and consolidating each tool overwhelms IT.

Cost surprises

Usage-based LLM pricing and overlapping licences make spend impossible to forecast.

WHY YOUR EXISTING STACK MISSES IT

The access happens **on the device.**

Network tools show you the traffic — which cloud AI services users hit, and how often. But it stops at the network edge. Even when the AI runs in the cloud, the access happens on the Mac: which tools are installed, what processes they spawn, what files they touch. None of that shows up in a DNS log.

CASB

Sees domains — not how an AI tool is configured or what it does on the device.

EDR

Watches for malware — an approved agent acting as the user looks normal.

Proxy

Inspects traffic — can't reach a local agent's config, MCP servers or commands.

THE CONFIGURATION PLANE FOR ENTERPRISE AI ON MAC

SEE IT

Full visibility

Every AI app, the commands it runs and the MCP servers it reaches — from first-party endpoint data no cross-platform tool can replicate.

CONFIGURE IT

Correct by default

A plain-language policy builder, enforced at the OS level over DDM. Tamper-resistant; applies even when tools route through third-party providers.

PROVE IT

Evidence on demand

A board-ready posture report and full audit trail, with a crosswalk to the EU AI Act — generated, not assembled.

Turn on visibility, and this is what you finally see — every AI tool, command and connection.

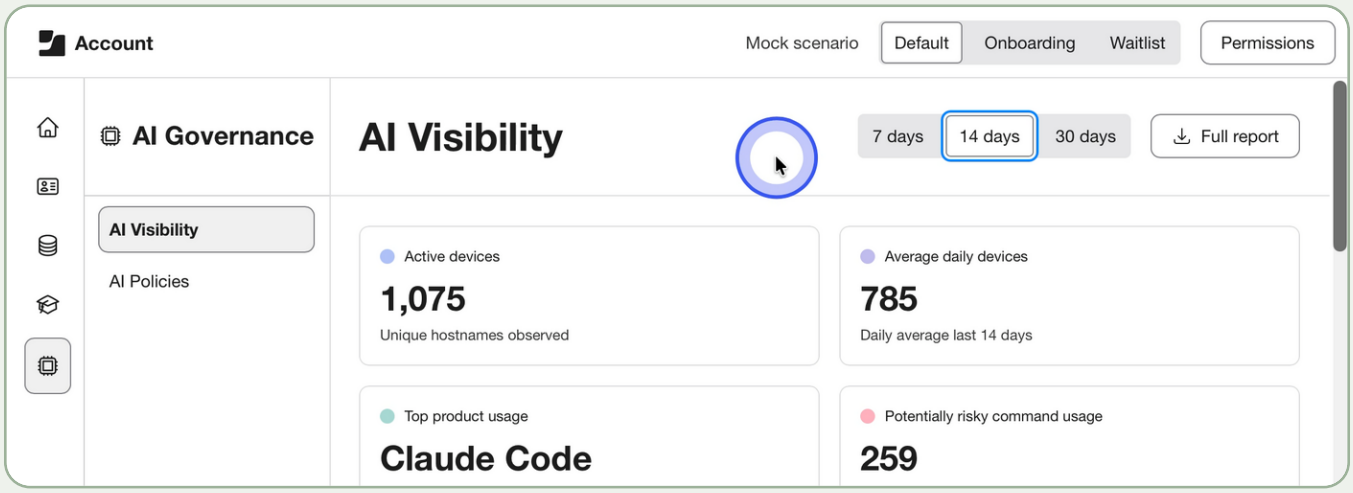

KEY CAPABILITIES

One enforcement plane, **six capabilities**.

AI application inventory

Surface every AI tool on the fleet — what launched it, what it connects to, and which invocations carry the highest risk.

MCP server inventory

See every MCP server running, what it exposes and which AI clients are connected — before you write a single rule.

Plain-language policy builder

Configure Claude Code, Claude Desktop and OpenAI Codex without editing JSON. Each setting labelled Enforced or Default.

OS-level, tamper-resistant enforcement

Policies deploy as Jamf Blueprints and apply at the OS level — users can't override enforced settings, and there's no provider bypass.

Vendor control-tracking engine

Monitors vendor releases and flags new or changed controls within days — policy stays current without a manual review cycle.

Governance report & audit trail

A board-ready Executive AI Posture Report: sanctioned vs ungoverned, full audit trail, EU AI Act crosswalk — generated automatically.

Plain-language configuration — every control a vendor exposes, without touching a schema.

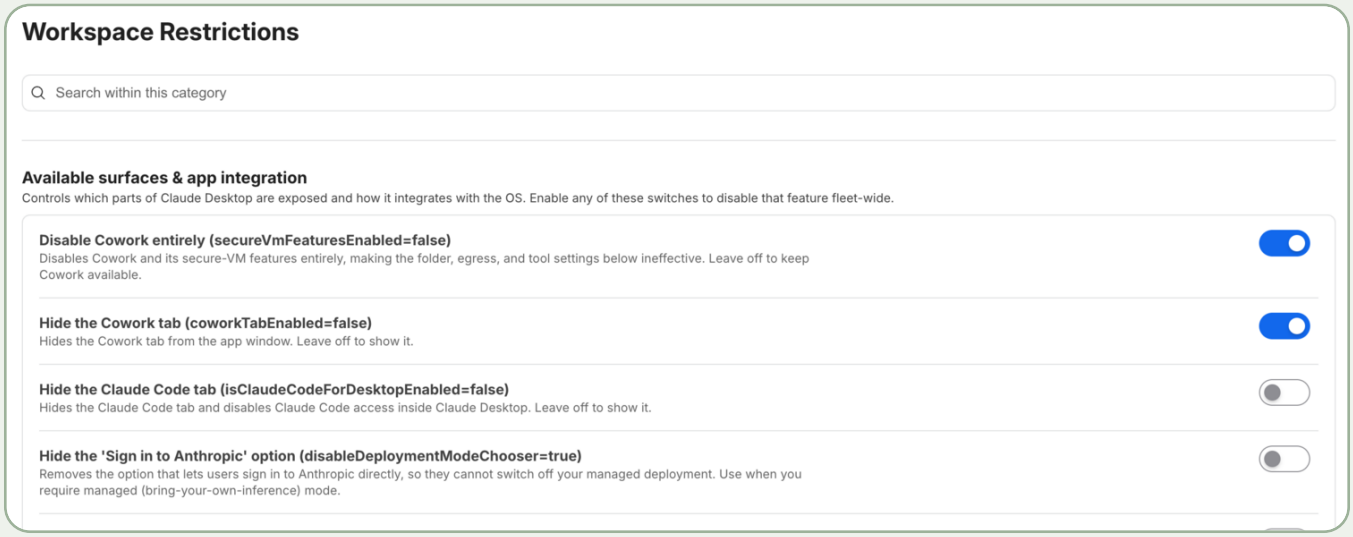

HOW IT WORKS

Included in the platform, deployed **the way everything else is.**

AI Governance isn't a bolt-on. Two flows run across three Jamf products you already operate — what you deploy and what you observe share state. No new agent, no new infrastructure.

OBSERVE — see what AI is running

DEVICE TELEMETRY
Jamf Protect agent
Monitors AI tool launches, MCP logs and commands run.

→

STORAGE
Reporting database
AI-specific events in a Jamf-hosted database.

→

API
Visibility API
Dashboards and reports in JSON.

→

JAMF ACCOUNT
Dashboard & report
Board-ready PDFs and drill-downs.

CONFIGURE — enforce what you've decided

JAMF ACCOUNT
Policy builder
Define the tool's managed configuration.

→

API
Policy API
Versioned policy definition.

→

JAMF PRO
Blueprint
Scoped and delivered over DDM.

→

DEVICE
Managed settings
Enforced at the OS level, no bypass.

THE FOUR PRINCIPLES IT'S BUILT ON

1. Gain visibility
You can't govern what you can't see — detect tools, processes and file access at the runtime.

2. Govern the tool
Guidance isn't enforcement. Set policy in the tool's own controls, not on the user.

3. Build it into deployment
Governance accompanies rollout, it doesn't react to it.

4. Use Apple-native tooling
The access happens on the device — a DNS log never sees it.

START WITH THE FOUNDATION

You can't govern what **you don't manage**.

AI Governance rides on the same platform you use to run your Macs. So the foundation for governing AI at scale is the foundation for managing Apple at scale — **Jamf for Mac**: management, security, identity and networking in a single per-device plan, built only for Apple.

Manage

Zero-touch enrolment, app lifecycle and patching, and managed OS updates — delivered declaratively.

Identity

Platform SSO, cloud identity sign-in, just-in-time accounts and passwordless workflows.

Protect

Purpose-built macOS security — threat prevention, CIS benchmarks and compliance visibility.

Connect

Zero Trust Network Access that follows the user — no legacy VPN friction.

AI Governance is included in Jamf for Mac. No new licence, agent or infrastructure. **Already a Jamf customer? You're closer to a defensible AI posture than you think.**

THE NEXT STEP

Request an AI Governance **Readiness Review.**

A short, no-obligation review with Onsite to see exactly what AI is running in your environment — and a plan to govern it, on the platform you already run. If you're a Jamf for Mac customer, we can switch it on; if you're not yet, we'll get the foundation in place first.

WHAT THE REVIEW INCLUDES

1. **Map your fleet against the prerequisites**
   We check where you are with Jamf for Mac, Jamf Pro and Jamf Protect, and what's needed to enable AI Governance.

2. **Switch on AI Visibility**
   Turn on the dashboard and see what AI is actually running across your Macs — apps, commands and MCP connections.

3. **Review the findings together**
   Walk through what's sanctioned, ungoverned and risky — and where your two exposure surfaces sit.

4. **Plan enablement**
   Agree the policy, rollout and the board-ready evidence you'll need — a clear path to a defensible AI posture.

Ready to see what's **really running** on your Macs?

Book your AI Governance Readiness Review with the Onsite team. Two weeks later, you're looking at your own dashboard — not our slides.

Request your Readiness Review → hello@onsiteit.africa

AI Governance is included in Jamf for Mac — no new licence, agent or infrastructure.

Prerequisites: a Jamf for Mac / Business · Jamf Pro and Jamf Protect in the same environment · Jamf Account SSO · the Jamf Protect agent deployed. Most existing Jamf customers already meet three of these.