ANTS Platform | AGENTIC ITAM

DOCUMENT TYPE EBOOK

PREPARED FOR CIOS, CISOS & IT LEADERS

EBOOK

The AI Challenges Defining IT Leaders' Success

Insights from 500 enterprise IT leaders on the five challenges of governing, funding and proving enterprise AI.

Contents

Executive summary		3
Challenge 1	“Our governance can’t keep pace with adoption.”	4
Challenge 2	“I don’t know what’s happening in my own organization.”	5
Challenge 3	“I’m being held accountable for systems I can’t control.”	6
Challenge 4	“Our CEO and board are misaligned, and external pressure is making it worse.”	7
Challenge 5	“I’m desperate for observability and auditability.”	8
What this means	Agentic ITAM: a system of record for enterprise AI	9
Governance AI	Guardrails for AI in production	10
Five questions to answer before the board asks		11
Methodology and sources		12

ABOUT THE DATA

The five challenges come from ANTS Platform’s research with 500 enterprise IT leaders. Supporting figures come from published research by Microsoft and LinkedIn, Gartner, IBM, SailPoint, CloudZero and MIT NANDA, each named where it is used. The methodology and every source are on page 12.

Executive summary

IT leaders are being asked to scale AI and answer for it at the same time, with controls built for an earlier era. The risk isn't AI itself. It's AI nobody can see, attribute or account for.

AI didn't enter the enterprise through procurement. It came in through browser tabs, desktop apps, coding assistants and API keys pasted into scripts, chosen by employees, adopted team by team and paid for on corporate cards and cloud bills. Agents followed, with credentials of their own.

The result is an AI estate most IT organizations can't fully see. Leadership asks how widely AI is used. Finance asks what it costs. The board asks what it returned. Auditors ask who approved it. Every one of those questions lands on the CIO, and few CIOs have a system that answers them.

Published research tells a consistent story. Employees bring their own AI tools to work. Security leaders suspect prohibited AI use inside their walls. Shadow AI already shows up in breach data, and it makes breaches costlier. AI agents take actions nobody intended. Spend climbs while confidence in its return lags.

None of this argues for slowing down. It argues for treating AI the way enterprises learned to treat every earlier technology wave: as assets that need an inventory, an owner, a budget and a record of value. This ebook sets out the five challenges we heard most often from 500 enterprise IT leaders, the published data behind each, and what changes when AI has a system of record. The methodology is on page 12.

95%

of organizations in MIT NANDA's study saw no measurable P&L return from their GenAI initiatives.

Measurement alone won't close that gap, but nobody closes it without measurement. To tell the AI that works from the AI that doesn't, an organization has to connect three things it usually keeps apart: who is using which AI, what it costs, and what it produces.

CHALLENGE 1

“Our governance can’t keep pace with adoption.”

For IT leaders, AI adoption now starts outside IT. Employees choose tools before IT reviews them, agents arrive inside SaaS products already under contract, and policy is always a step behind.

Employees aren’t waiting

In Microsoft and LinkedIn’s 2024 Work Trend Index, 78% of AI users said they bring their own AI tools to work, and the pattern held across every age group. The approved-tools list an IT leader signs off on is a partial view at best.

Policy is still being written

IBM’s 2025 Cost of a Data Breach Report found that 63% of breached organizations had no AI governance policy or were still developing one.

Their security peers see it too

In a Gartner survey of 302 cybersecurity leaders, 69% said they suspect or have evidence that employees use prohibited public GenAI. Gartner predicts that by 2030 more than 40% of enterprises will have security or compliance incidents linked to unauthorized shadow AI.

Why blocking doesn’t hold

Block lists chase thousands of AI tools, many embedded in apps you already allow. Blocking pushes usage onto personal devices, where it is even harder to see.

78%

of AI users bring their own AI tools to work

69%

of cybersecurity leaders suspect or have evidence of prohibited GenAI use

63%

of breached organizations had no AI governance policy, or were still writing one

HOW ANTS HELPS **Discover AI**

Governance that keeps pace starts with a live inventory: AI tools and agents across browsers, desktop apps, CLIs and IDEs, each tied to a named user, department and device through Okta or Azure AD, and classified by risk so review effort goes where exposure is highest.

CHALLENGE 2

“I don’t know what’s happening in my own organization.”

IT leaders are expected to know what runs in their estate. With AI, many can’t say how many tools and agents are in use, who uses them or what they cost.

Invisible to existing tools

AI activity rarely reaches SIEM, CASB or endpoint logs with a person attached. A request to a model provider looks like ordinary HTTPS traffic, and an API key in a script has no owner field.

Agents multiply quietly

In SailPoint’s survey of 353 IT and security professionals, 82% of companies use AI agents but only 44% have governance policies for them. Agents often hold several machine identities of their own.

Spend nobody can attribute

In CloudZero’s 2025 survey, 57% of organizations still track AI costs in spreadsheets and 15% have no formal tracking at all.

Ownership is the starting point

When every tool, agent and key has a named owner, approval and accountability stop being arguments after the fact.

AI agents: adoption runs ahead of governance



Source: SailPoint / Dimensional Research, 2025

HOW ANTS HELPS **Identity attribution · Budget AI keys**

Every AI event is tied to a named person, department and device. Apps and agents get one revocable ANTS key each, with an owner and tags, while provider credentials stay held centrally and are never handed out.

CHALLENGE 3

“I’m being held accountable for systems I can’t control.”

When an agent acts on its own or an employee’s AI tool leaks data, the questions are who, what and on whose authority. Without control or records, the answer defaults to the IT leader.

Agents act on their own

80% of organizations in SailPoint’s survey said their AI agents had performed unintended actions, including accessing unauthorized systems (39%), sharing inappropriate data (33%) and downloading sensitive content (32%).

Credentials leak

23% said their agents had been tricked into revealing access credentials.

Shadow AI raises the bill

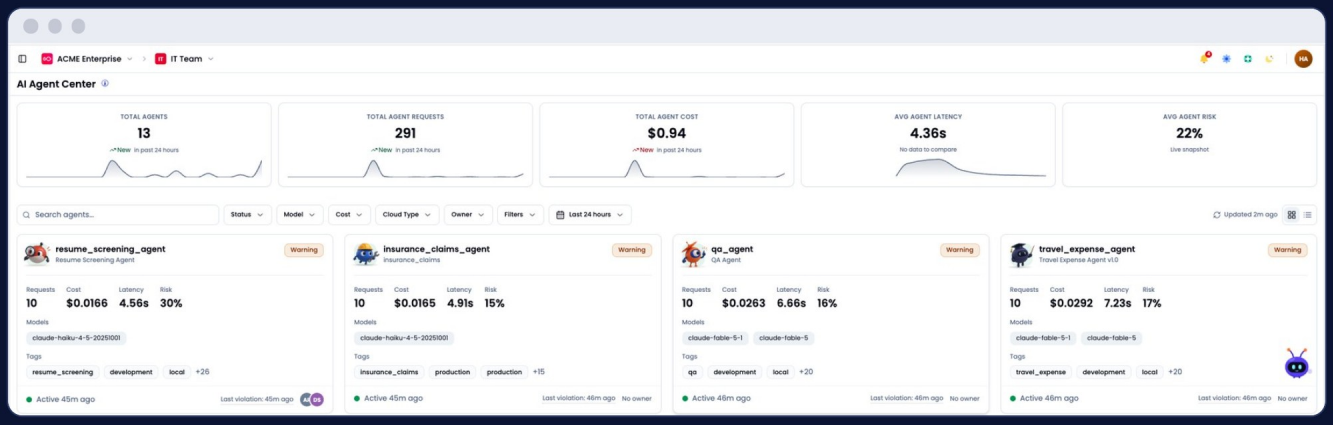
In IBM’s 2025 study, one in five organizations reported a breach involving shadow AI. High levels of shadow AI added an average of $670,000 to breach costs.

Access controls were missing

Of organizations that had an AI-related breach, 97% lacked proper AI access controls.

HOW ANTS HELPS **AI Agent Center · Governance AI**

Control that matches the accountability: every agent with its cost, risk, owner and last violation; Governance AI guardrails that catch tool misuse and privilege violations at runtime; and hard budgets that refuse requests at the limit.



AI Agent Center: agents with cost, latency, risk score, models, owner and last violation. Demo tenant.

CHALLENGE 4

“Our CEO and board are misaligned, and external pressure is making it worse.”

Headlines and vendors tell CEOs to move faster on AI. Boards ask what the spend has returned. IT leaders sit in between, asked for speed and proof at the same time.

Spend is rising fast

Average monthly AI spend in CloudZero’s 2025 survey was $85,521, up 36% from 2024.

Proof is scarce

MIT NANDA found that 95% of organizations saw no measurable P&L impact from GenAI, despite an estimated $30–40 billion in enterprise investment.

Projects get cut

Gartner predicts more than 40% of agentic AI projects will be canceled by the end of 2027 over costs, unclear value or weak risk controls.

Alignment needs one set of numbers

Only 51% of organizations in CloudZero’s survey were confident they could evaluate the ROI of their AI costs.

95%

of organizations saw no measurable P&L impact from GenAI

36%

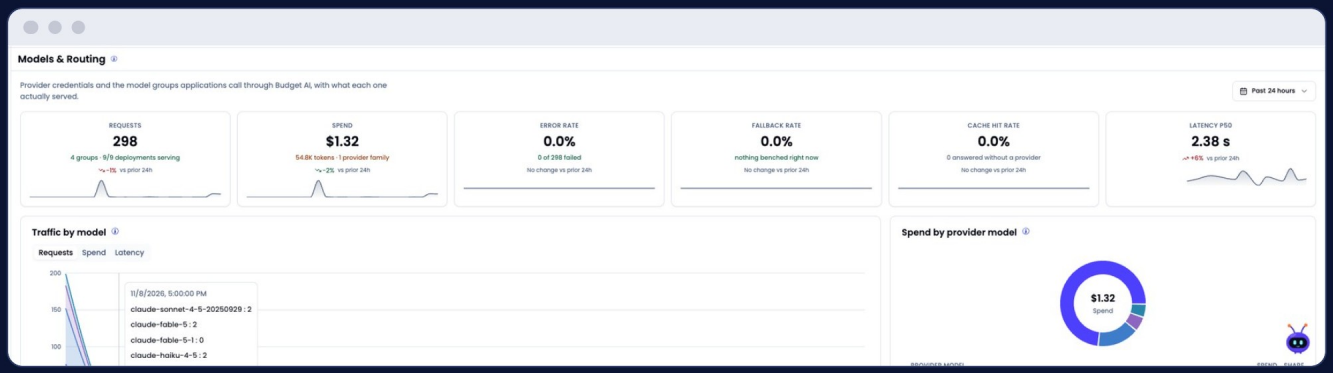
year-over-year rise in average monthly AI spend

51%

are confident they can evaluate the ROI of their AI costs

HOW ANTS HELPS **Budget AI · Gauge AI**

One set of numbers for the CEO, CFO and board: spend by key, department and model as requests happen, and Gauge AI usage and behaviour showing where AI is working and where it isn’t.



Budget AI · Models & Routing: requests, spend and latency by model. Demo tenant.

CHALLENGE 5

“I’m desperate for observability and auditability.”

IT leaders want to see what AI is doing across their systems, trace any action back to a person, and show an auditor the record without a scramble.

Audits aren’t happening

IBM found that among organizations with AI governance policies, only 34% regularly audit for unsanctioned AI, and just 22% run adversarial testing on their AI models.

Frameworks expect a record

NIST AI RMF and ISO/IEC 42001 expect an inventory of AI systems, named owners and records of use.

Gartner’s advice

Gartner recommends that CIOs set enterprise-wide AI usage policies, run regular audits of shadow AI activity and include GenAI risk in SaaS assessments.

What observability means for AI

Every request tied to a key, an owner, a model and a cost, kept as a searchable record.

34%

of organizations with AI policies regularly audit for unsanctioned AI

22%

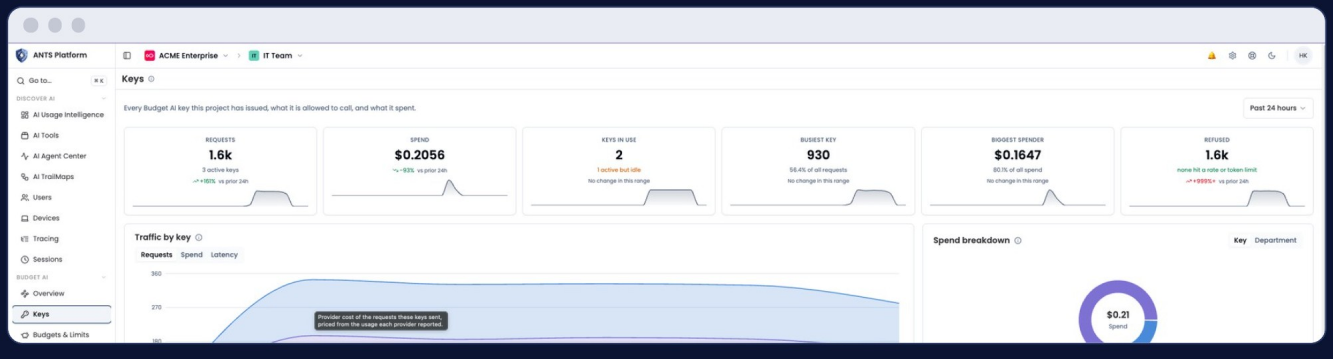
run adversarial testing on their AI models

37%

of breached organizations had AI approval or oversight processes

HOW ANTS HELPS **Audit & oversight · Governance AI**

Audit logs, request tracing, SIEM export and Governance AI drift detection. Every key shows what it called, what it spent and who issued it. The Endpoint AI Agent and browser extension collect metadata only, never prompt content.



WHAT THIS MEANS

Agentic ITAM: one record for enterprise AI

Servers, software licenses, SaaS and cloud each eventually got a system of record: an inventory with owners, costs and a lifecycle. AI is the next asset class, and it moves faster than any of them.

The five challenges share a root cause. AI is in use without an inventory, without owners, without budgets that hold and without a record of what it produced. Fixing them one at a time with point tools leaves the seams where problems hide. ANTS Platform is built as one record across all four.

1 Discover

DISCOVER AI

Inventory every AI tool, agent, model and key in use, each attributed to a person and device.

2 Measure

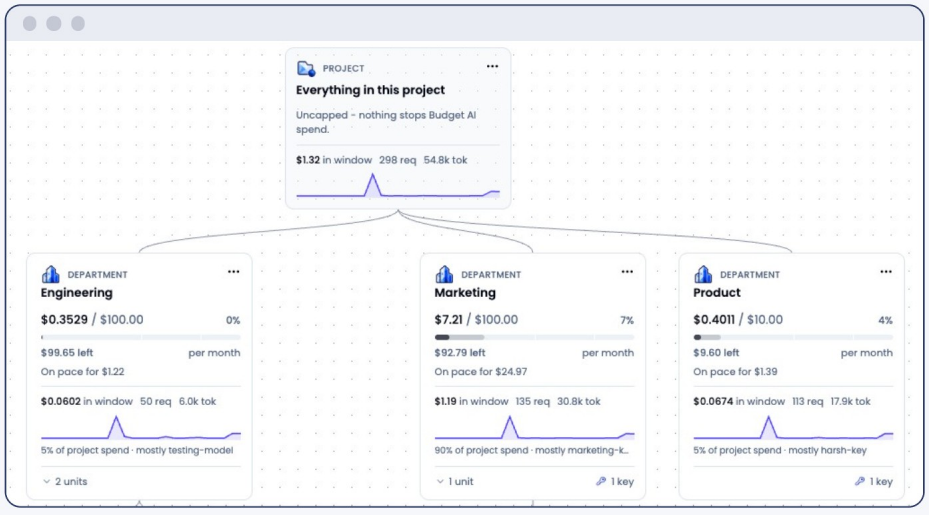
BUDGET AI · GAUGE AI

Cost by key, team and org unit, budgets that enforce, and usage that shows where AI works.

3 Govern

GOVERNANCE AI · RECORDS

Runtime guardrails, and every event and costed request recorded with its owner for the board and the auditor.



BUDGET AI · BUDGETS & LIMITS

Budgets that mirror the org chart

Project, department and sub-department budgets, each with spend against its limit, what is left this month and where it is on pace to land. Units without their own limit draw on their parent.

Demo tenant.

1 record

per AI tool, agent, model and key, tied to a named person and device

0 prompts

collected by the Endpoint AI Agent or browser extension: metadata only

Hard

budgets refuse requests at the limit rather than only warning

GOVERNANCE AI

Guardrails for AI in production

Visibility tells IT leaders what AI is doing. Governance AI decides what it is allowed to do, watching agents and AI applications at runtime and enforcing policy before a violation reaches a customer, a regulator or the board.

Drift detection

Spots when a model's or agent's behaviour, outputs or risk shift away from its baseline over time.

Tool misuse

Flags agents calling tools they shouldn't, or using approved tools in unapproved ways.

Privilege violations

Catches agents and users acting beyond the access and permissions they were granted.

PII, PCI and PHI

Detects personal, payment and health data in prompts and responses before it leaves your control.

Ethics and bias

Checks AI outputs for biased or unethical content, with a record of every flagged decision.

Custom guardrails

Your own input and output rules, written for your policies, your industry and your regulators.

Every violation is logged with its full trace context and can trigger an action: log, alert, block or redact.

CHALLENGE 3

Control over the agents and AI apps IT leaders are held accountable for.

CHALLENGE 5

An auditable record of every policy decision, ready for the security team and the auditor.

YOUR AI PROGRAM

Evidence that frameworks such as NIST AI RMF and ISO/IEC 42001 ask organizations to keep.

Five questions to answer before the board asks

If your team can answer these from a system rather than a spreadsheet, the five challenges become manageable.

01	**How many AI tools and agents are in use, and by whom?**	DISCOVER AI Live inventory with identity attribution and risk classification.
02	**Who owns each one, and was it approved?**	KEYS · TRUSTED MODELS One revocable key per app or agent, with owner and tags; a list of approved models.
03	**What are we spending, by team and by app?**	BUDGET AI Spend by key, project and org unit; hard and soft budgets.
04	**Where is AI actually working?**	GAUGE AI Usage patterns and behaviour, effectiveness scoring, automation opportunities.
05	**Can we show an auditor the record?**	GOVERNANCE AI · AUDIT Guardrail violations, audit logs, tracing and SIEM export, every event tied to its owner.

HOW TO START

STEP 1

Mutual NDA

ANTS provisions an isolated tenant, with SSO through Okta or Azure AD.

STEP 2

Security review

Your infosec and architecture teams review data handling and deployment.

STEP 3

Scoped pilot

Agree success criteria and get a first AI inventory within the pilot window.

Methodology and sources

ABOUT THE DATA

Percentages quoted in this ebook come from the published studies below, not from the ANTS research. Each describes its own sample and they are not directly comparable. Where a figure is a forecast, the text says so.

SUPPORTING RESEARCH

1. Microsoft and LinkedIn, *2024 Work Trend Index: AI at Work Is Here. Now Comes the Hard Part*, May 2024.
   Survey of 31,000 people in 31 countries.

2. Gartner, "Gartner Identifies Critical GenAI Blind Spots That CIOs Must Urgently Address," press release, November 2025.
   Survey of 302 cybersecurity leaders, March–May 2025; 2030 figure is a Gartner prediction.

3. IBM, *Cost of a Data Breach Report 2025*, July 2025.
   Research conducted by Ponemon Institute across breached organizations worldwide.

4. SailPoint, *AI Agents: The New Attack Surface*, May 2025.
   Global survey of 353 security and IT professionals, conducted by Dimensional Research.

5. CloudZero, *The State of AI Costs in 2025*, May 2025.
   Survey of 500 US software professionals, manager level and above, at organizations of 250–10,000 employees, March 2025.

6. MIT NANDA, *The GenAI Divide: State of AI in Business 2025*, July 2025.
   153 leader surveys, 52 interviews and analysis of 300+ public AI implementations.

7. Gartner, "Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027," press release, June 25, 2025.
   Gartner prediction.

ANTS Platform

Agentic ITAM.

Find out what AI is running in your enterprise, who owns it, what it costs and what it returns. Start with a scoped pilot and a first AI inventory.

Book a scoped pilot · antsplatform.ai

ABOUT ANTS PLATFORM

ANTS Platform is the system of record for enterprise AI: one identity-attributed record of every AI tool, agent, model and key, with the controls and evidence to govern it. Discover AI, Budget AI, Gauge AI and Governance AI share one data plane. Product screens shown are from a demo tenant.