

THE MOVING AI **FRONTIER**

THE 2026 GUIDE TO AI TRAINING
AND ENABLEMENT AT ENTERPRISE SCALE

**How to Turn Your Workforce's AI Momentum into
Structured Capability and Lasting Competitive Advantage.**

Contents

Foreword - Dan Newton, CEO, Security Journey.....	3
Executive Summary - The Commercial Stakes, in Plain Terms	4
Part One: The State Of AI Adoption	
Chapter 1: The Workforce That Didn't Wait.....	5
Chapter 2: What Shadow AI Actually Costs.....	8
Part Two: The Chasm	
Chapter 3: Why Most Enterprise AI Programs Don't Work.....	10
Chapter 4: The Governance Gap Nobody Wants to Own.....	12
Part Three: Building the Bridge	
Chapter 5: Celebrate the Citizen Builders	13
Chapter 6: The Three Layers That Actually Compound.....	15
Chapter 7: The Regulated Industry Reality	16
Part Four: The Secure Enablement Audit	17
Part Five: The Moving Frontier	
Chapter 8: The Frontier Doesn't Stand Still	18
Conclusion	21

Dan Newton

CEO, Security Journey

As I write this in June of 2026, TechCrunch published a piece this week titled 'Everyone is navigating AI security in real time — even Google,' it stopped me. Not because it was surprising, but because it named exactly what we are seeing every day across the enterprises we work with. The transition period described at Google Cloud - where security can't be bolted on after the fact and can't be left to employees to figure out alone — is the operating reality for every organization with more than a few hundred people. The difference is that most of them don't have Google's resources to work through it.

This guide was written because the conversation our team has been having with CISOs, CTOs, and AI leaders is not a theoretical one. It is a conversation about what is already happening — the finance analyst who connected ChatGPT to the ERP, the creative team building synthetic media with tools nobody reviewed, the operations lead who automated a workflow in a weekend without telling IT. These people are not reckless. They are resourceful. The question we keep getting asked is not how to stop them. It is how to build the organizational infrastructure that lets them keep going safely — and how to extend that capability to the thousands of employees who are watching from the sidelines waiting for someone to show them how.

The answer this guide offers is not a technology solution. It is a structural one. Governance that travels. Enablement that is role-specific and hands-on. Measurement that gives boards something real to look at. The organizations that build this now will not just avoid the incident. They will have a workforce that moves faster on every AI capability that arrives next year and the year after. That compounding advantage is what this guide is about.



The Commercial Stakes, in Plain Terms

The business cost of unstructured AI adoption is now measurable. A shadow AI breach adds an average of \$670,000 to an organization's incident cost. (IBM, 2025) 42% of executives say the process of adopting generative AI is tearing their company apart. (Writer, 2025) And organizations without structured AI enablement are losing ground on three fronts simultaneously: security exposure from ungoverned tools, sales friction as prospects ask harder questions about AI governance, and competitive position as AI-native peers widen the capability gap with every model release.

Inside every organization, three distinct groups have emerged. Early adopters are accelerating ahead — building workflows, connecting tools, and shipping AI-assisted work without waiting for permission or policy. Dabblers — the largest group — see the opportunity but lack the role-specific direction to translate curiosity into consistent capability. Laggards remain hesitant, often not out of resistance but out of uncertainty about where to start. The organizations crossing the chasm are the ones that have built the structures to bring all three groups forward together.

For many organizations, where that structure is missing, the result is shadow AI: employees using tools and building solutions outside approved systems because the official infrastructure has not kept pace with demand. This is not reckless behavior. It is a rational response to an enablement vacuum. The consequence is a widening gap between what employees are doing with AI and what their organizations are equipped to support — one where workflows remain fragmented, capability remains uneven, and competitive advantage slips to those who move faster.

This guide names what is happening, what is failing, what is working, and what Secure AI Enablement looks like when built to compound rather than decay. It draws on direct conversations with CTOs, CISOs, and AI leads navigating this in real time, and on research from McKinsey, Gartner, Forrester, IBM, and Verizon.

88%



of organizations say they use AI in at least one business function.

McKinsey, 2025

63%



of organizations lack AI governance policies despite widespread adoption.

IBM, 2025

1%



have mature AI deployments delivering real business value. That gap is the chasm.

McKinsey via Ant Murphy, 2025

Ch.1 THE WORKFORCE THAT DIDN'T WAIT

Your People Moved First. Now Your Organization Has to Catch Up.

While your AI task force was still debating acceptable use policy, your finance team connected ChatGPT to Ramp, Pigment, and your ERP — systems that contain vendor contracts, headcount costs, board-level financials, and customer spend data — because it made their month-end close faster. There is no data processing agreement governing that flow. No audit trail. No retention policy you approved.

Your marketing team fed unpublished campaign briefs, competitive positioning documents, and customer segmentation models into Claude and GPT-4o to write sharper copy faster. Your creative team went further: cloning executive voices in ElevenLabs, generating product demo videos in HeyGen, producing synthetic media in Runway — using real brand assets, real faces, real scripts. Your operations lead automated three vendor workflows over a weekend without IT involvement. Your sales team assembled Clay, Apollo, and Gong into an AI layer sitting across customer call recordings almost certainly governed by GDPR — without a single procurement request.

These people are not reckless. They are resourceful. And as Francis de Souza, COO of Google Cloud, observed: this shadow AI behavior — employees reaching for consumer tools without organizational oversight — is not unique to any one company. It is the default operating state of enterprise AI in 2026.

→
88% of organizations say they use AI. 1% have deployments delivering real business value. That 87-point gap is not a technology problem. It is a structural one — and it is the gap this guide exists to close.


“Security is not something you can bolt on later, and it’s not something you can leave up to employees to do on their own.”





Francis de Souza
COO, Google Cloud
TechCrunch, May 2026

THE 8/72/20 – WHERE YOUR WORKFORCE ACTUALLY SITS

CISO Grant Kahn of Versapay, who has built AI enablement programs inside real organizations, describes the workforce distribution he consistently encounters: 8 percent are self-starters who will build and experiment regardless of what the organization does. They don't need enablement — they need guardrails. Twenty percent won't engage meaningfully regardless of how good the program is. And the 72 percent in the middle are where everything is decided. They need structure, role-specific context, and evidence that someone has thought about their actual job — not AI in the abstract. Grant's phrase: 'the ROI battleground.'

8% 
Early Adopters
Build Regardless

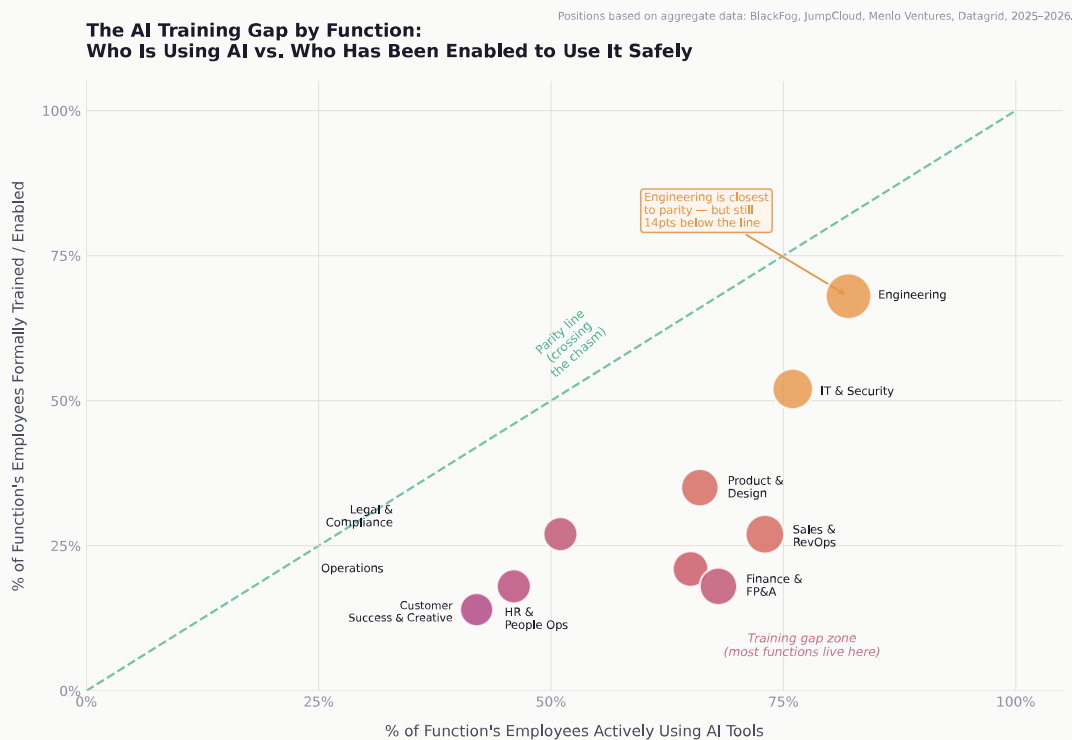
72% 
Dabblers
Willing, often curious, but need structure and role-specific context to unlock.

20% 
Laggards
Regardless of Design
 McKinsey via
 Ant Murphy, 2025

Source: Grant Kahn, CISO, Versapay — direct workforce enablement experience, May 2026. Mirrors Rogers' Diffusion of Innovations (1962).

THE TRAINING GAP BY FUNCTION

The below scatter maps the same pattern at a functional level — how much of each team is using AI versus how much has been formally enabled. Every bubble sits below and to the right of the parity line. Some dramatically so.



This pattern is not new. Everett Rogers mapped it in 1962 in Diffusion of Innovations — the same adoption curve that governed electricity, the telephone, the internet, and the smartphone. Innovators and early adopters move fast. The early and late majority — the dabblers — need to see peers succeeding before they commit. Laggards come last or not at all. Geoffrey Moore identified the gap between early adopters and the early majority as so significant it warranted its own name. What Rogers and Moore documented at a societal level, Grant Kahn is watching play out inside individual organizations every day.

The implication is direct: the 8 percent are already building. The 20 percent aren't your problem. The 72 percent in the middle — the dabblers who are curious but directionless — are where an AI enablement program either creates compounding value or quietly fails.

“The intern, or the 75-year-old executive whose grandkids showed him vibe coding. They don't know enough to know enough. And they're very confident in their ignorance.”

Data Scientist at a Global Consulting and Technology Firm March 2026

THE DATA BEHIND THE DISTRIBUTION

98%

of organizations have employees using unsanctioned AI apps — whether leadership knows it or not.

BlackFog / Programs.com, 2026

49%

of employees at 500+ person companies use AI without employer approval — including 69% of C-suite members.

BlackFog, 2026

32%

of employees have received any formal AI training at work. The other 68% are self-teaching on company data.

SQ Magazine, 2026

78%

of AI users in the enterprise bring their own tools to work — unsanctioned, unmonitored, unmanaged.

Microsoft Work Trend Index, 2024

48%

of employees want formal AI training — they see it as the clearest path to improving adoption.

McKinsey, 2025

40%

fewer security incidents at companies with formal AI training programs vs. those without.

SQ Magazine, 2026

The Risk Isn't That They're Using AI. It's That You Don't Know How.

Every function has assembled its own AI stack. None of it went through procurement. None of it has been reviewed for data handling. Most of it is invisible to the people responsible for keeping the organization safe. The following table is not a hypothetical risk map — it is a description of what is running in most 500+ person organizations right now.

Function	Tools Already in Use	The Data Risk — With Numbers
Finance & FP&A	ChatGPT, Claude, Pigment, Ramp AI, Notion AI	Headcount costs, vendor contract terms, board-level financials, and customer spend data flowing through models with no data processing agreement. 23% of employees admit to inputting company financial information into unsanctioned AI tools. (BlackFog, 2026)
Marketing & Creative	HeyGen, Synthesia, ElevenLabs, Runway, Midjourney, Sora	Executive likenesses cloned for content. Unreleased campaign briefs and competitive positioning used as generation inputs. Real brand assets feeding models whose training terms changed. IP ownership is genuinely unclear.
Sales & RevOps	Clay, Apollo AI, Gong, Lavender, ChatGPT	Customer call recordings, contact data, and deal intelligence connected through AI layers. 86% of employees surveyed use AI weekly for sales tasks. (BlackFog, 2026) Almost certainly governed by GDPR or CCPA.
Operations & Product	Notion AI, Make, Zapier AI, ChatGPT, v0	Internal process documentation, vendor work flows, and product road maps fed into personal accounts. Citizen builders pushing AI-generated outputs into production without review — silently breaking compliance definitions of who counts as a 'developer'.
Legal & Compliance	Harvey, Spellbook, ChatGPT, Notion AI	Legal AI active integration nearly doubled from 14% to 26% in 2025 alone. (Datagrid) Contract language, M&A terms, and employment disputes run through models with inconsistent enterprise-grade data handling.
HR & People Ops	Paradox, Eightfold, ChatGPT, Beamery	27% of employees admit to revealing employee data (salary, performance) to unsanctioned AI tools. (BlackFog, 2026) AI-assisted candidate screening creating employment liability most employment lawyers haven't fully mapped.

The CISO who was once the person saying no is now the person setting the course. That transition — from blocker to director — is one of the clearest signals that shadow AI has moved from a technical concern to a leadership one. As de Souza put it this week: **“This is a board-level issue and an executive team issue. It’s not just a security team’s issue.”**

The Data Exposure by Category

(BlackFog, 2026):

33% of employees share enterprise research or datasets with unsanctioned AI tools

27% reveal employee data (salary, performance) to AI tools outside company oversight

23% input company financial information into unmanaged models

None of this requires malicious intent. It requires an enablement vacuum.



Access Is Not Enablement. Here's What the Difference Costs.

Gartner has found that two-thirds of organizations have not yet scaled AI beyond pilots or experimentation, and only a small fraction are achieving meaningful enterprise-wide results. The question is not why organizations are slow to adopt AI — their employees have already answered that question independently. The question is why the programs designed to guide that adoption keep failing in the same five ways.

“Most leaders are mistaking basic access or adoption metrics for transformation. This ‘enablement illusion’ is hiding risks and draining ROI.”



Swagatam Basu
Senior Director
Analyst, Gartner



FAILURE MODE 1 // Procurement Is Not a Strategy

Buying Anthropic's Claude or Microsoft's Copilot is not a strategy. It's a procurement event.

The rollout announcement goes out. Licenses are provisioned. A blog post goes live about the company's commitment to AI. Ninety days later, active usage is below 30 percent. The power users are the same people who were building with AI before the license existed, and everyone else has an icon on their desktop they've clicked twice. Access is not enablement. The Copilot Adoption Hub is not a curriculum. Giving 5,000 employees access to an AI tool without the structured capability layer is the enterprise equivalent of buying everyone a gym membership and calling it a fitness program.

FAILURE MODE 2 // Developers Trained. Everyone Else Forgotten.

The technical teams got structured training. The other 10,000 got a newsletter.

Coding accounts for 55% of all departmental AI spend in 2025. (Menlo Ventures) The result: one function in a 12,000-person organization has sandboxed environments, structured curricula, and security support. The analysts, PMs, operations leads, creative teams, finance, legal, HR — the other 10,000 — got a quarterly lunch-and-learn and a Slack channel with resources nobody reads. Adaptavist's research found that AI training is predominantly offered to high earners, meaning AI fluency is concentrating at the top of the income distribution and creating a two-tier workforce. That is not a rollout. It's a pilot left running while the rest of the org builds its own shadow stack.

FAILURE MODE 3 // Generic Training Teaches Nothing Specific

Generic content is expensive, forgettable content.

A 90-minute mandatory course that gestures at AI without addressing anyone's actual workflow is not enablement. The finance team's AI fluency needs look nothing like the legal team's. The creative director building synthetic video content is navigating completely different risk terrain than the account executive using AI for deal support. When training is designed for everyone in the abstract, it works for no one in particular. The completion rate looks fine. The capability delta closes by approximately zero.

FAILURE MODE 4 // A PDF Policy Is Not Governance

An acceptable use policy that arrives as a PDF is not governance.

56% of workers say they lack clear guidance on AI usage policies. (SQ Magazine, 2026) Governance that lives in a document travels nowhere. Real governance is embedded in workflow, measurable, visible to the people it's supposed to guide, and owned by someone with genuine accountability. Most enterprise AI policies were written under time pressure, formatted as dense PDFs, distributed via email, and never seen again. The analyst in finance opens ChatGPT on Monday morning and does exactly what they were doing before — because nothing in the policy told them what to actually do differently.

FAILURE MODE 5 // No Owner. No Measurement. No Outcome.

If nobody owns it, nobody can improve it.

42% of executives say the process of adopting generative AI is tearing their company apart. (Writer, 2025) AI enablement in most large organizations sits at the intersection of IT, HR, Legal, and the Office of the CTO — which means it's owned by everybody conceptually and nobody practically. No single accountable leader. No leading indicators being tracked. No feedback loop between what the program delivers and what the business needs. The steering committee meets quarterly. The program decays between meetings.

The Org Chart Problem Sitting at the Center of Enterprise AI

There is a pattern that shows up inside almost every enterprise navigating AI at scale: the program is technically everyone's responsibility and practically nobody's. Technology leadership owns the infrastructure. HR owns the learning program. Legal owns the policy. The CISO owns the risk framework. Each builds its piece in isolation. The pieces don't connect. And the employee sits in the middle of a confusing patchwork of competing guidance, waiting for something coherent that never quite arrives.

“There's no such thing as an AI strategy without a data strategy and a security strategy. They need to go hand in hand.”



Francis de Souza
COO, Google Cloud
TechCrunch, May 2026

OWNERSHIP FAILURE

The Committee Is Not an Owner

A steering committee that meets monthly cannot move at the speed AI adoption is moving. By the time a decision is socialized, approved, and communicated, the tools have updated, the risk surface has shifted, and the employees who were waiting for guidance have stopped waiting.

BOARD PRESSURE

Audit Committees Are Starting to Ask

Boards are asking not 'do we have an AI policy?' but 'can you demonstrate that AI risk is managed?' Those are different questions. A document is an acknowledgment. Telemetry, ownership, and outcomes are evidence. The organizations that can answer the second question have a significant advantage.

THE CISO SHIFT

From Blocking AI to Governing It

CISOs are increasingly placed on AI Governance Councils with enablement mandates alongside their security mandates. The role has shifted from blocker to director. Organizations that haven't made this structural move are running a security posture designed for a different problem.

THE BUDGET SIGNAL

Whoever Claims This Gets the Budget

AI enablement budget is real and growing. In many organizations it currently sits unallocated because no single leader has claimed it with a credible program behind the claim. The first functional leader who walks in with a measurable plan attached to business outcomes tends to get the budget.

Ch.5 CELEBRATE THE CITIZEN BUILDERS

The People Already Doing This Are Your Most Valuable Asset.

The operations lead who rebuilt their monthly vendor workflow in GPT-4o without asking IT. The finance analyst who automated their reconciliation process on a Saturday. The account manager who cut their pre-call research from 45 minutes to 8 using a custom Clay prompt. These people have crossed the chasm personally. Most organizations are treating them as a compliance problem.

Andrew Ng has argued consistently that the barrier to AI adoption inside organizations is not technical capability — it is organizational permission. When employees self-enable at this level, they are not demonstrating recklessness. They are demonstrating that the will, the aptitude, and the job-specific motivation are already present. The missing ingredient is not ambition. It is structure.

STEP ONE | Find Them. Surface Them.

Run a shadow AI audit as a talent identification exercise, not a disciplinary one. Who has built something? What did it unlock? These are your AI champions — currently operating without organizational recognition or structural support.

STEP TWO | Let Them Teach the Middle

The workflows citizen builders create are often the best source material for role-specific training. When the ops lead who automated invoice processing teaches the rest of the team — in their language, with their tools — the 72% in the middle finally has something that applies to them.

STEP THREE | Draw the Line Between Innovation and Exposure

Give citizen builders a framework that lets them keep building without creating compliance exposure they're not equipped to see. The goal is not to slow them down. It is to make their momentum durable and safe to scale.

STEP FOUR | Build the Champion Network

The 72% in the middle don't need a motivational speaker. They need a peer who has done it, in a role like theirs, with tools like theirs. Champion networks — structured programs that formally recognize and resource internal AI leaders — are the highest-leverage investment most organizations aren't making.



CASE STUDY

monday.com's Champions Layer

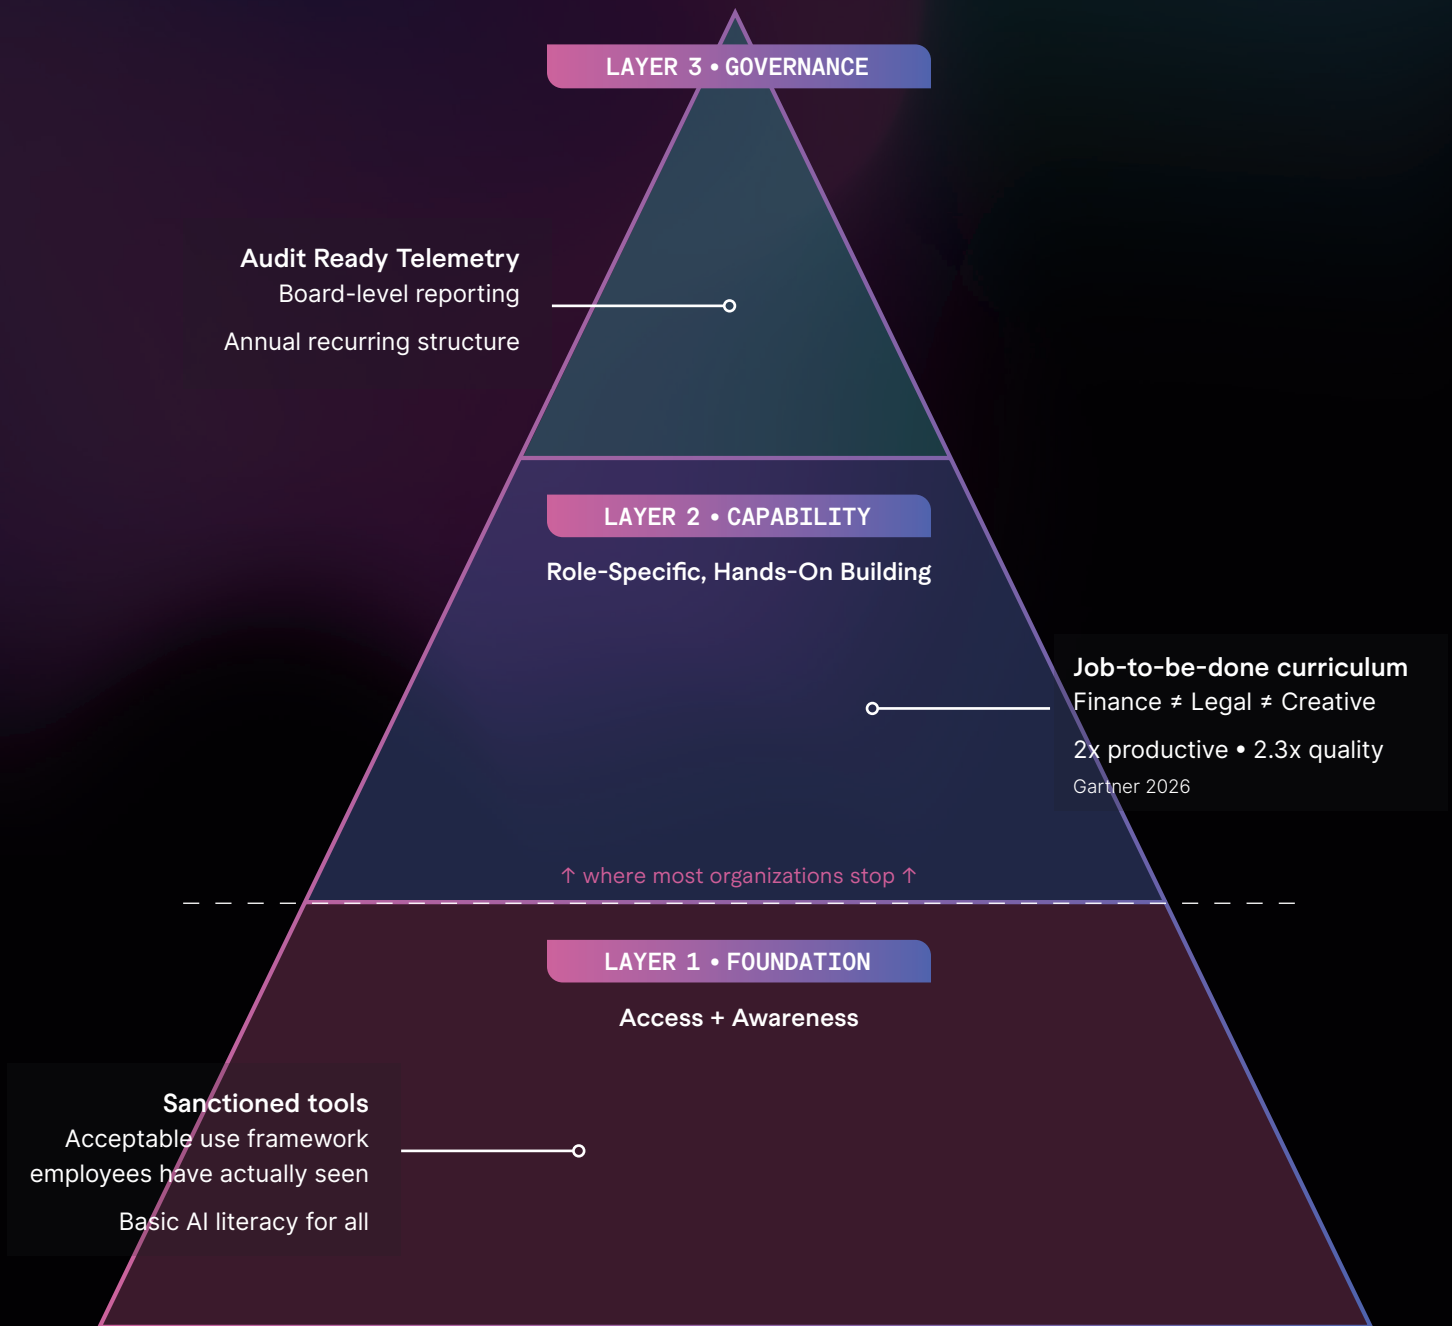
monday.com built what they call the Champions Layer: a structured internal program where employees from every department — marketing, HR, finance, operations — are formally resourced to build their own AI-powered work flows using tools like n8n, without engineering dependency.

The cultural shift they were engineering: from 'AI is something built for me' to 'AI is something I build myself.' That transition — from passive recipient to active builder — is what separates organizations that cross the chasm from those that stay on the wrong side of it.

Source: Keren Koshman, Internal AI Innovation Leader, monday.com · Medium, November 2025.

The Organizations Crossing the Chasm Don't Have More Tools. They Have More Structure.

The Forrester finding is precise: it is not the tools that are missing. It is the structure around them. AI enablement infrastructure that lasts is built in three layers. Skip to Layer 3 without Layer 1 and 2 and you get governance without capability. Stop at Layer 1 and you get access without behavior change. Only all three together compound — each one enabling the next.



skip a layer and the whole thing collapses – each one enables the next

For Financial Services, Healthcare & Media: The Stakes Are Different

Every enterprise is navigating the same underlying dynamic. But in regulated industries, the blast radius of getting it wrong is not a learning moment — it is a material event. The EU AI Act, emerging US frameworks, and sector-specific regulators are creating compliance pressure that most general industry guidance hasn't priced in. The smarter organizations in regulated sectors are not treating this as a burden. They are using it as a budget-unlocking mechanism.

INDUSTRY Financial Services	INDUSTRY Healthcare	INDUSTRY Media & Entertainment
<p>Model risk management frameworks (SR 11-7 in the US) were not designed for generative AI — but regulators are extending them to cover it. Financial services leads all sectors with over 50% of tech budget going toward AI. (IBM/Kore.ai, 2025) An analyst using an unsanctioned model to support a credit decision creates model risk exposure with no validation process. Customer data under GDPR, trading data with market sensitivity, M&A information that may constitute MNPI — all potentially flowing through unmanaged channels.</p>	<p>71% of nonfederal acute care hospitals report using predictive AI — but only 22% have it properly governed. (Datagrid) HIPAA has no 'we didn't realize it was in the prompt' exception. AI tools influencing care pathways face increasing FDA scrutiny. As LinkedIn CISO Lea Kissner told the New York Times this week: 'We're going to need people to deal with the bug-pocalypse' — and she doesn't expect the industry to understand AI security in any sustainable way for at least several years.</p>	<p>The creative AI stack — HeyGen, Synthesia, ElevenLabs, Runway — is moving faster in media than almost any other sector, and the legal frameworks are genuinely unsettled. IP ownership of AI-generated content, right of publicity issues with cloned voices and likenesses, training data questions — all active litigation areas with no clear precedent. Insurance for AI-related IP disputes is still being priced by underwriters who don't fully understand the exposure.</p>
<p>Governance Trigger: Compliance mandate + vendor rationalization pressure = a conversation that starts with the CISO or CRO, not IT.</p>	<p>Governance Trigger: HIPAA audit question, a near-miss with patient data in a prompt, or a new AI-assisted clinical workflow proposed by operations.</p>	<p>Governance Trigger: A talent contract renewal with AI terms, a deepfake incident involving company assets, or a rights dispute over AI-generated content.</p>

When a non-developer pushes AI-generated code to production, the compliance definition of 'developer' breaks. Most organizations don't realize it's happened until an auditor asks the question."

Dan Newton
CEO, Security Journey

Five Questions. If You Cannot Answer Three with Evidence, You Are Exposed.

Most organizations believe they have an AI governance posture. Most are wrong. A breach tied to an unsanctioned AI tool averages \$670,000 in added cost. AI governance gaps create sales friction: enterprise buyers now ask about it in procurement reviews. And board-level scrutiny of AI risk is no longer optional. Five questions. Answer them honestly, with evidence.

- 1 Do you know which AI tools your employees are actually using — sanctioned and unsanctioned?**
Not the tools you approved. The tools they're using. Including the ones connected to Ramp, Pigment, Gong, and the ERP. 98% of organizations have employees using unsanctioned apps — but most can't enumerate them. If you can't map your shadow AI stack, you can't govern it.
- 2 Does your acceptable use policy exist in a format employees have actually encountered and understood?**
A PDF in a SharePoint folder is evidence of a policy, not governance. 56% of workers say they lack clear guidance on AI usage policies. (SQ Magazine) If your employees couldn't summarize your policy in two sentences, it isn't reaching them.
- 3 Do you have role-specific AI capability programs in place — or just org-wide access?**
The finance team, the legal team, the creative team, and the operations team have completely different AI jobs-to-be-done. If your enablement program doesn't address that, it's addressing nothing specifically and therefore nothing effectively.
- 4 Who owns AI enablement outcomes in your organization — and what are they measured on?**
A steering committee is not an owner. A named leader with a clear mandate, a budget, and measurable outcomes they're accountable for is an owner. If you can't name that person immediately, the program doesn't have one.
- 5 Can you produce board-ready evidence that AI risk is managed — not just acknowledged?**
Your board is starting to ask this. 'We have a policy' is an acknowledgment. Audit-ready telemetry showing which tools are in use, which populations are trained, and what the governance layer covers is evidence. These are fundamentally different things.

All five with confidence and evidence:	Your organization is materially further ahead than most. The work now is scaling what you've built and connecting the layers.
Two or three:	You have foundations but significant gaps. Fault lines are likely showing up somewhere in your program. Structured enablement is your next step.
One or zero:	You have access, not enablement. The chasm is open. The organizations that move fastest in the next twelve months will set the standard their industries follow.

Ch. 8 THE FRONTIER DOESN'T STAND STILL

Crossing the Chasm Gets You to the Frontier. The Frontier Keeps Moving.

The organizations in the top-right quadrant of the chart at the front of this guide are not waiting for the frontier to arrive. They are already operating at it. What follows is not a prediction. It is a description of what the leading edge looks like right now — and what will be unremarkable across high-performing organizations within six months.

The implications for enterprise AI enablement are significant and immediate. Here is what is already normalizing at the front of the market:

THE WORKSPACE SHIFT

Work Happens Inside the Agent Now

The model of 'AI inside your SaaS tool' is already being displaced. The next model: employees bring their own AI agent and use their SaaS tools inside it. The agent sees everything, holds all context, and operates across applications. Enablement programs built only around tool-specific training will be obsolete within two product cycles.

NON-TECHNICAL ROLES WIN

PMs and Designers Are Shipping Faster Than Anyone

The people winning the AI era right now are not the best coders. They are the best thinkers. Product managers with deep user empathy and spiky product sense are shipping faster than engineering-heavy teams. Full-stack designers are building their own visions without handoffs. The citizen builder is becoming the most valuable person in the room.

THE SUPER-AGENT MODEL

One Company-Wide Agent. One Owner. Trickles Down.

The early assumption — every employee has their own personal AI agent — is already being revised. The model that is actually working: one forward-deployed, AI-savvy person maintains a single company-wide agent (think Shopify's River). Specialized team agents follow as the infrastructure matures. This is an ownership and governance question as much as a technology one. Shipper's own company doubled in size this year despite being deeply AI-forward — because every automation still needs a human making sure it keeps working.



SPOT LIGHT

Dan Shipper

CEO of Every — one of the most-watched AI-native operators in the market right now — describes his current working reality:

He spends all his time inside Codex. Writing documents, managing email, doing research — everything. He uses Google Docs, PostHog, and every other tool he needs within the agent's in-app browser. The agent can see what he's doing and has all of his context. He and his agent work together. This is not a demo. This is Tuesday.

The governance implication of all of this is direct. An AI enablement infrastructure built for the current tool landscape — individual SaaS applications, discrete use cases, function-level training — will need to extend into agent-level governance within six months. Who owns the company-wide agent? What data does it access? What is the audit trail when it acts on behalf of an employee? These are not hypothetical questions. They are questions Shopify, Every, and a growing number of organizations are answering right now, at the frontier.

The organizations that have built the three-layer model in this guide are the ones positioned to absorb this shift without starting over. Their governance layer can extend. Their champion network can test and translate. Their measurement infrastructure can adapt. The organizations still debating whether to write an acceptable use policy are not just behind the current moment. They are two transitions behind the frontier — and the frontier is not waiting.

COMPETITIVE SIGNAL

AI Fluency as a Talent Differentiator

In tech and media especially, AI capability is becoming a visible part of employer brand. The organization that can credibly tell candidates ‘here is how we invest in your AI fluency’ — and back it up with a real program — is competing differently for the exact people who will matter most in the next eighteen months.

BOARD CONFIDENCE

A Track Record, Not Just a Policy

Organizations that built audit-ready AI governance in 2024 and 2025 are walking into 2026 board meetings with something competitors don’t have: demonstrated, measured, improving governance over time. As agent-level AI introduces new board-level questions, that track record is the foundation every answer will be built on.



CONCLUSION

Three Moves. This Quarter.

Three moves. This quarter. Execute them.

Move 1

Run the Secure Enablement Audit. Answer the five questions in Part Four honestly, with evidence. Not with what your policy says. With what your employees are actually doing, what your governance actually covers, and who actually owns the outcome. This is your baseline.

Move 2

Identify Your 72%. Find the dabblers in your organization — by function, by role, by team. These are the people who are willing but directionless. They are your ROI. Map what they need that they are not getting. That gap is your program.

Move 3

Assign Ownership. Name the person who owns Secure AI Enablement outcomes in your organization. Give them a mandate, a budget, and a measurement framework. A steering committee is not an owner. A named leader with accountability is. Everything else follows from this.

THE FRONTIER KEEPS MOVING. THE ORGANIZATIONS THAT MAKE THESE THREE MOVES THIS QUARTER WILL BE THE ONES SETTING THE STANDARD — NOT CATCHING UP TO IT.

“The time for individual experimentation has passed.”



Grant Kahn
CISO, Versapay
May 2026

AI enablement that upgrades human skill and organizational *velocity*

See how Security Journey's AI Advantage platform can get your entire organization AI-ready — governed, enabled, and ahead of the curve.

go.securityjourney.com ›

SOURCES & ACKNOWLEDGEMENTS

Customer and prospect quotes are drawn from direct conversations conducted by the Security Journey team between February and May 2026, used with permission. A full permissions pass is recommended before publication.

- [1] McKinsey, *The State of AI in 2025: Agents, Innovation, and Transformation*, 2025
- [2] Ant Murphy, *The AI Chasm*, antmurphy.me, 2025
- [3] Verizon, *2026 Data Breach Investigations Report*, 2026
- [4] Gigamon, *Hybrid Cloud Security Survey*, 2026
- [5] IBM, *Cost of a Data Breach*, 2025
- [6] BlackFog / Programs.com, *Shadow AI Survey*, 2026
- [7] Reco AI, *State of Shadow AI*, 2025
- [8] Microsoft, *Work Trend Index*, 2024
- [9] McKinsey, *Superagency in the Workplace*, 2025
- [10] SQ Magazine, *AI Adoption and Training Survey*, 2026
- [11] Gartner, *Predicts by 2027: 50% of Enterprises Without a People-Centric AI Strategy Will Lose Their Top AI Talent*, 2026
- [12] Forrester, *The State of AI*, 2025
- [13] Adaptavist, *Crossing the AI Chasm: How CTOs Can Drive Adoption*, 2025
- [14] Menlo Ventures, *State of Enterprise AI*, 2025
- [15] Datagrid, *Legal AI Adoption Report*, 2025
- [16] IBM / Kore.ai, *Enterprise AI Investment Report*, 2025
- [17] JumpCloud, *IT Trends Report*, 2026
- [18] Writer, *Enterprise AI Report*, 2025
- [19] Keren Koshman, *Crossing the AI Chasm: Investing in People First*, Medium, 2025
- [20] Connie Loizos, *Everyone is navigating AI security in real time — even Google*, TechCrunch, May 24 2026
- [21] Thomas Otter, *Workday and a16z, Work in Progress (Substack)*, April 2026
- [22] Everett Rogers, *Diffusion of Innovations*, Free Press, 1962
- [23] Geoffrey Moore, *Crossing the Chasm*, HarperBusiness, 1991
- [24] Dan Shipper, CEO, *Every — observations on AI-native working*, LinkedIn / X, May 2026