

# AI Adoption: The Complete Enterprise Workbook (2026)

*The definitive guide to understanding, identifying, and encouraging impactful AI adoption within your organization.*

---

## Measuring and Encouraging AI Adoption

Shortcuts:

- [The high cost of enterprise AI tools](#)
- [Scaling up AI, from pilots to company-wide deployment](#)
- [Building useful AI dashboards](#)
- [Adoption benchmarks by industry](#)
- [Adoption benchmarks by company department](#)
- [Driving adoption in non-technical teams](#)

In [AI Adoption: The Complete Enterprise Guide](#), we described AI-powered tools and what it takes to successfully adopt them.

In this Workbook, we show you how to actually find AI tools in your workplace; compare your findings to industry benchmarks; and encourage AI adoption within your organization.

Successful AI adoption proceeds in two directions:

- **From the bottom up.** Individuals find tools that help them do their work. It's very common for users to create a ChatGPT subscription for personal use, whether paid or free; the software has nearly 1 billion weekly active users at this writing. Users can also easily slide from Google Search into Google Gemini, which delivers "no-click" answers to search queries and offers additional functionality that rivals ChatGPT.
- **From the top down.** Companies reach subscription deals or hosting arrangements for AI-powered software: LLMs such as ChatGPT, Claude, and Gemini; AI-first tools such as Cursor for code or ElevenLabs for audio; conventional software augmented with AI feature sets, such as Notion AI, Slack AI, and Microsoft Copilot; vertical solutions designed for specific industries; and in-house projects that use AI to varying degrees.

Each type of adoption has its own benefits and concerns. We want to point these issues out here, rather than describe them. So we've created a brief table summarizing the advantages and disadvantages of each.

	Bottom-up	Top-down
Initial idea	User	Company
License holder	User	Company
User enthusiasm	Tends to be high	Varies
Typical use pattern	Few licenses, used more intensely	Many licenses, often with many unused
Security challenges	Many, serious (including data leakage)	Some, less serious (data leakage minimized)
What success looks like	Individual success, limited w/o company adoption	Training, encouragement, rewards lead to high usage
Visibility to management	Licenses and use invisible without special software	Licenses visible; poor visibility into usage
Management concerns	Security (high concern), governance, low adoption	Security (moderate concern), governance, low adoption

Table 1. Some advantages and disadvantages of user-led vs. company-managed AI adoption (without an AI measurement and management platform)

While every individual and company situation differs, these descriptions reflect industry reports and our experience from talking to customers.

How many AI tools should your people be using? Our report, [The State of Enterprise AI 2026](#), gives some numbers that you might consider as a guideline.

We asked our survey respondents two broadly related questions:

- Do you expect your company to receive ROI from AI in the next 6 months?
- How many AI tools do you use?

As shown in Figure 1, we found that those with low or no expectation of achieving ROI used an average of 1.1 AI tools. While those with a high expectation of achieving ROI used an average of 2.7 tools, nearly three times as many.

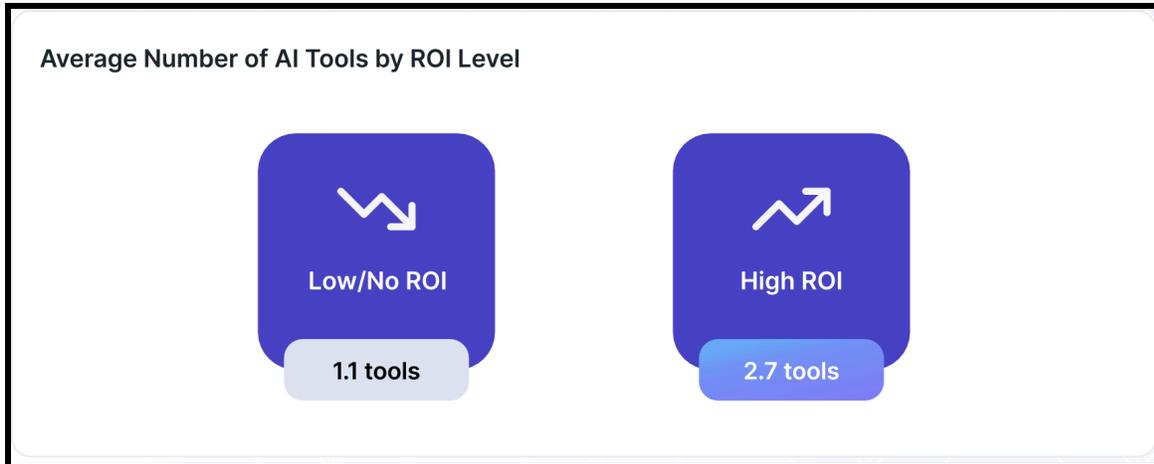


Figure 1: Expectation of achieving ROI with AI vs. how many AI tools the respondent uses (from [The State of Enterprise AI 2026](#))

We interpret this to mean that higher AI achievers, in higher AI-achieving organizations, tend to use several tools, whereas those with less AI engagement may only use one.

There are, of course, exceptions to every rule:

- A company that focuses its entire AI effort on getting its large software development staff to use Anthropic's Claude, for example, while providing training and sharing best practices, might be well on its way to achieving ROI for this single deployment
- Whereas a company that reimburses employees for any AI software tool they care to try, with little measurement or follow-through, might have many mini-deployments, with little hope of achieving (or even knowing if it has achieved) ROI

It's early in the growth and adoption of AI, but our survey may give you a directional indication as to what you want to achieve in your organization.

---

## Calculating the (High) Cost of Enterprise AI Tools

Most enterprises can tell you what they spend on AI licenses. Almost none can tell you what AI actually costs.

### The Iceberg: License Fees Are Just the Tip

When the CFO asks "what are we spending on AI?", the answer is almost always the procurement number—per-seat license fees, neatly tracked and easily totaled up. That number is real. It is also wildly incomplete.

**License and subscription fees represent only 20-35% of total AI implementation costs.** Gartner's research confirms this range. The remaining 65-80% of costs are hiding in infrastructure, integration, training, and governance costs, shadow AI spend, and the productivity drag of employees using tools poorly. Organizations routinely underestimate total AI cost by **40-60%**, with budget overruns of 30-40% in the first year.

Think of it as an iceberg. Above the waterline: subscription line items. Below: six categories of cost that are real, recurring, and almost never aggregated. When organizations run full discovery, the true number is consistently **three to five times** what the procurement dashboard shows.

---

## The Seven Categories of Costs

A complete **AI tool TCO assessment** accounts for seven distinct categories—but most organizations only track the first one.

### 1. License and Subscription Fees

**Microsoft 365 Copilot** runs **\$30 per user per month**. **GitHub Copilot Enterprise** costs **\$39 per user per month**. **ChatGPT Enterprise** typically lands around **\$60 per user per month**. These are three tools. Most enterprises run dozens.

For a 10,000-person enterprise deploying Copilot broadly, license fees alone clear **\$3.6 million annually**—before a single integration is performed or training session conducted. And that is one tool in a portfolio that typically contains **200-300 AI tools**, once you count what IT does not officially track.

### 2. Infrastructure and Compute Costs

API calls, cloud compute, GPU allocation, data storage, and scaling costs that grow non-linearly as usage expands: **65% of IT leaders report unexpected charges from consumption-based AI pricing**, with actual costs exceeding estimates by 30-50%. Organizations building custom models or retrieval-augmented generation (RAG) pipelines face compute costs that can dwarf their license spend.

### 3. Integration and Customization

Connecting AI tools to CRM, ERP, HRIS, and data warehouses is where “we will just plug it in” meets reality. **Organizations underestimate integration costs by 30-50%**. Legacy system connections require 25-35% more investment than projected. Large enterprises regularly spend more than \$150,000 per major tool integration.

### 4. Training and Change Management

**BCG's 10-20-70 rule** holds: 70% of AI value comes from people, processes, and data—not algorithms or technology. Organizations should budget **\$2,000-5,000 per**

**employee** for comprehensive AI upskilling. [OpenAI's 2025 State of Enterprise AI report](#) documents a **6x productivity gap** between power users and average employees using the same tools. Closing that gap requires continuous **proficiency development**, not a one-time session. Organizations that allocate 15-20% of total AI budget to training consistently outperform those that minimize training.

### 5. Governance and Compliance Overhead

Management tasks, including monitoring usage, enforcing data policies, maintaining audit trails, and responding to regulatory requirements, demand dedicated resources. [IBM's 2025 Cost of a Data Breach report](#) found that breaches involving AI cost **\$4.63 million on average**. EU AI Act fines reach 35 million euros. A single compliance incident can erase years of productivity gains. Our [AI Governance Checklist](#) provides the framework.

### 6. Unsanctioned ("Shadow") AI Costs

This category produces the three-to-five-times multiplier. As detailed in our [Governance Guide](#), **98% of organizations have employees using unsanctioned AI tools** and **76% have active bring-your-own-AI usage**. CIOs estimate that their employees are using 60-70 AI tools; actual monitoring reveals 200-300. [IBM's 2025 Cost of a Data Breach report](#) found that shadow AI breaches cost **\$670,000 more** than average incidents, and one in five organizations has already experienced a breach tied to shadow AI.

### 7. Opportunity Cost of Low Proficiency (The "AI Tax")

This may be the largest cost category—and it is the one most TCO calculations miss entirely. [Workday's January 2026 study](#) found that **37% of time "saved" by AI is lost to rework**. Only **14% of employees** consistently achieve net-positive outcomes. This is the **AI Tax**: the hidden productivity drag when people use AI when they lack the proficiency to use it well.

If Copilot saves the average employee 45 minutes per day, the AI Tax means roughly 17 minutes are consumed by rework. Across 10,000 employees, that is approximately **2,833 hours of rework per day**—invisible on dashboards, but tangible in missed deadlines and inconsistent quality. The AI Tax is not set in stone. It is a function of proficiency, and it shrinks when organizations invest in closing the 6x productivity gap between power users and average employees using the same tools.

---

## What a 10,000-Person Enterprise Actually Spends

Cost Category	Estimate	Notes
<b>License fees</b>	\$5-8M	M365 Copilot alone costs \$3.6M for 10K licenses
<b>Infrastructure and compute</b>	\$2-4M	API calls, cloud, storage

Cost Category	Estimate	Notes
<b>Integration and customization</b>	\$1-3M	CRM, ERP, HRIS connections
<b>Training and change management</b>	\$2-5M	\$200-500/employee annually, ongoing
<b>Governance and compliance</b>	\$1-2M	Staff, tooling, audit, legal
<b>Shadow AI spend</b>	\$2-5M	Unsanctioned tools, duplicates
<b>Rework / AI Tax</b>	\$3-8M	37% rework across thousands
<b>Total estimated TCO</b>	<b>\$16-35M</b>	vs. \$5-8M in tracked licenses

The license line represents **roughly 25% of true cost**. [KPMG's Q4 2025 AI Pulse Survey](#) found enterprises project deploying **\$124 million on AI** annually. If the hidden cost multiplier applies, real exposure is substantially higher than any budget document captures.

---

## The TCO Calculation Framework

Build your cost picture at three levels:

**Per-tool TCO** = License cost + allocated infrastructure + shadow duplicates + training + integration + governance + rework cost. This is what procurement needs for vendor negotiations.

**Per-department TCO** = Sum of tool TCOs + department change management + compliance costs + rework burden. This is what department heads need for budgeting.

**Per-workflow TCO** = Tool costs allocated to the workflow + integration + workflow-specific rework. This is the most actionable view—connecting cost directly to outcomes. The [AI Impact Guide](#) provides the framework for tying these costs to value creation.

---

## When the Math Works—and When It Doesn't

**High adoption + high proficiency = ROI.** License cost amortizes across real productivity gains. Rework stays low. The Copilot ROI Framework details how to measure this state.

**Low adoption = expensive shelfware.** S&P Global reports that **42% of companies** have abandoned most AI projects. A tool with 10,000 licenses and 2,000 active users has an effective per-user cost **five times the sticker price**.

**High adoption + low proficiency = the AI Tax.** This scenario looks productive on dashboards but bleeds value. People use the tools; outputs require constant rework. Net productivity is marginal or negative—invisible without quality measurement alongside adoption metrics.

The breakeven test: if a \$30/user/month tool does not save each user at least \$30/month in **net** productivity after rework, the investment is underwater.

---

## Cost Optimization Strategies

**Consolidate ruthlessly, experiment cheaply.** Run full AI tool discovery within your organization. Identify duplicate licenses and unsanctioned consumer licenses. Enterprise analysts project 2026 will see organizations spend more through **fewer vendors**. Our [AI Transformation Guide](#) details the process.

**Negotiate with usage data, not headcount.** The gap between “licenses purchased” and “licenses actively used” is your leverage. Negotiate license counts down to sensible levels. Explore usage-based pricing where vendors offer it.

**Invest in proficiency to shrink the AI Tax.** Reducing the rework rate from 37% to 20% across 10,000 employees is worth more than any license discount. Proficiency has the highest ROI of any cost optimization lever.

**Reassess quarterly.** The best tool today may be a commodity—or a feature of another tool—in six months. Favor shorter contract terms over multi-year agreements.

**Measure cost-per-outcome, not cost-per-seat.** The real question is not “what does a license cost?” but “what does this outcome cost with AI versus without it?” The [AI Transformation Guide](#) provides the metrics framework.

---

## Common Cost Mistakes

**Counting only license fees.** Without the six below-the-waterline categories, your TCO understates reality by three to five times.

**Ignoring shadow AI spend.** With 98% of organizations running unsanctioned tools, the untracked spend is not a rounding error—it is a budget category larger than most IT line items.

**Not measuring usage against spend.** An adoption problem is a cost problem. Fix adoption before negotiating license discounts.

**Locking into multi-year contracts.** Microsoft's Copilot pricing has already shifted toward promotional rates of \$18/user/month in early 2026—organizations locked into older agreements missed the adjustment.

**Treating training as a one-time expense.** The 6x productivity gap does not close just by scheduling a webinar. Budget for ongoing employee development.

**Ignoring cost avoidance.** Governance that prevents a \$4.63 million data breach creates enormous value that never appears in a cost-savings analysis.

---

*Larridin gives enterprises the visibility to calculate true AI TCO—discovering every tool in use, measuring adoption and proficiency by team and workflow, and connecting usage data to the financial metrics that separate productive investment from expensive shelfware. [See how Larridin quantifies your real AI cost ->](#)*

---

## How to Scale AI From Pilots to Enterprise-Wide Deployment

---

88% of AI pilots never reach production. Here is the five-phase framework that separates organizations that scale from those trapped in pilot purgatory.

### The Scaling Gap: Why 80%+ of AI Pilots Never Reach Production

The pilot worked. The demo impressed leadership. And then—nothing.

This is not an edge case. It is the default outcome. IDC found that for every **33 AI prototypes** a company builds, only **4 make it to production**. RAND documented an **80%+ AI project failure rate**, twice the rate of non-AI IT projects. Gartner predicted that **30% of generative AI projects would be abandoned after proof of concept** by end of 2025. McKinsey found that while 88% of organizations use AI, only **one-third have begun scaling** organization-wide. BCG found just **5% achieve value at scale**.

The industry has a name for this: **pilot purgatory**. Nearly two-thirds of organizations remain stuck there. The question is not whether AI works—it does. The question is what the organizations that escape pilot purgatory do differently.

What are you trying to achieve? Our report, [The State of Enterprise AI 2026](#), shows what the highest AI achievers get out of using AI, vs. “the rest of us.”

We asked our survey respondents for their time savings from using AI; how many hours/month they got back by using AI tools.

The distribution is a classic power-law breakdown; the top group averages more than four times the time savings of the bottom group. Figure 2 shows the breakdown.

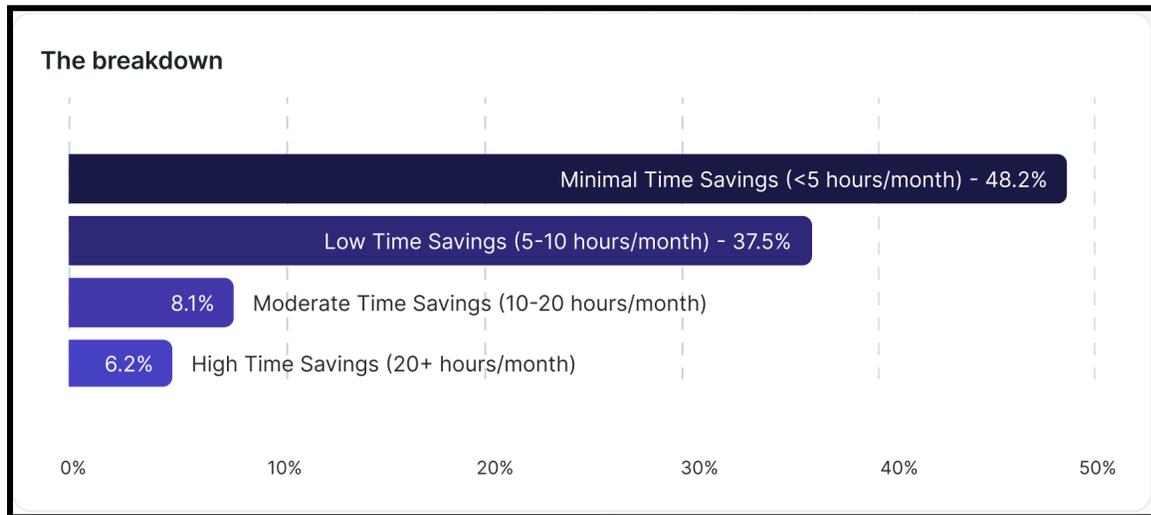


Figure 2: Time savings from using AI and the number of respondents achieving each level (from [The State of Enterprise AI 2026](#))

We interpret this to mean that diligent effort to learn AI, at least where suitable tools exist for an employee's major tasks, can pay off in truly impressive productivity increases.

---

## Three Failure Modes

### Technical: Infrastructure Wasn't Built to Scale

A pilot runs on a sandbox and a single API key. Enterprise deployment requires production-grade data pipelines, model monitoring, security controls, and system integration. [Kyndryl's 2025 Readiness Report](#) found that **62% of AI projects stall due to infrastructure gaps**. The pilot team solved a narrow problem with enthusiasm. Scaling requires architecture.

### Organizational: Change Management Was an Afterthought

[BCG's analysis](#): **70% of AI value creation** comes from people, processes, data, and measurement—not the technology. The [AI Transformation Guide](#) calls this the **inner orbit**. The pilot team was self-selected early adopters. The broader organization includes skeptics, the undertrained, and the resistant. [OpenAI's 2025 State of Enterprise AI report](#) documented a **6x productivity gap** between power users and typical employees. Scaling without addressing the human side is scaling with a 6x handicap.

## Strategic: The Wrong Problem Was Piloted

Many pilots start with “where can we use AI?” instead of “what business problem needs solving?” [RAND](#) identified this as a root cause: stakeholders often miscommunicate what problem needs solving. When the pilot cannot articulate measurable business impact, it cannot survive the budget conversation required for scaling.

---

## The Scaling Framework: Five Phases

### Phase 1: Validate

Before scaling, pressure-test whether the pilot deserves to scale. Can you name the business metric it improved? Can you quantify the improvement? The AI Maturity Model diagnostic questions apply: what is the primary metric you are optimizing, and what are your guardrail metrics?

Many pilots prove AI can perform a task. **Validation proves that performing the task creates measurable business value.** These are not the same thing. Build measurement from day one using the five impact dimensions from the [AI Impact Guide](#)—effectiveness, quality, time, revenue, and cost. This data is the foundation of every subsequent phase.

### Phase 2: Standardize

The phase most organizations skip—and why most scaling efforts fail—sounds simple: Convert a one-off success into a repeatable system.

**Infrastructure:** Move from ad-hoc data access to production-grade pipelines.

**Governance:** Establish policies, data handling rules, and the governance spectrum from the [AI Transformation Guide](#). Governance feels like overhead during a pilot; at scale, ungoverned AI is a compliance risk. **Monitoring:** Deploy telemetry using the [AI Adoption Guide](#) five-step program—Discover, Classify, Establish Metrics, Instrument, Report and Act. **Proficiency baseline:** Assess where the broader workforce sits on the proficiency spectrum. Identify Level 3-5 users from the pilot—they become your champions.

### Phase 3: Enable

Scaling a tool without scaling the capability to use it is the most common scaling mistake.

**Champion networks.** Deploy pilot team power users as embedded mentors in expansion teams. They carry contextual knowledge and credibility that IT-mandated rollouts never will. **Differentiated enablement.** Level 1-2 users need foundations. Level 3 users need advanced coaching. Level 4-5 users need orchestration training. **Change management.** [Zapier achieved 97% AI adoption](#) through culture—hackathons, experimentation time,

shared norms. [Meta tied AI impact to performance reviews](#). Both worked. What does not work is assuming that deploying a tool drives adoption.

Training can make a huge difference in AI effectiveness. For our report, [The State of Enterprise AI 2026](#), we asked our respondents whether their organization had one or more formal AI training programs. We also asked about their AI proficiency, user satisfaction with AI, and AI-powered productivity gains. Figure 3 shows the results.

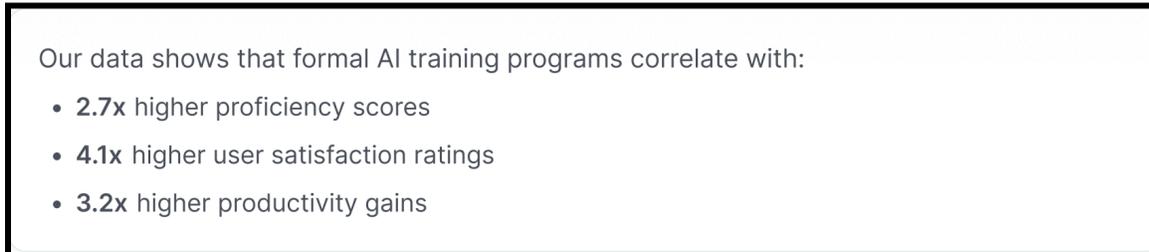


Figure 3: Formal AI training programs are associated with higher AI proficiency, user satisfaction, and productivity gains.  
(from [The State of Enterprise AI 2026](#))

We interpret this to mean that having a formal AI training program in an organization pays off. We are aware that “correlation does not equal causation”; just starting an AI training program, where there are headwinds against success with AI, may not vault the company up the AI effectiveness charts. But organizations making a strong structural effort, including formal AI training, seem likely to achieve results that are well ahead of the curve.

## Phase 4: Scale

Don't expand from one team to the entire organization in a single step. Scale in concentric rings.

Select expansion targets using Phase 2 data—highest-value workflows, greatest **Capacity Reallocation Value**, strongest readiness. Each expansion follows the **90-day Transformation Loop** from the AI Transformation Roadmap: Assess, Prioritize, Execute, Measure, Adapt. Each ring follows the same discipline: new baseline, new measurement cycle, champion support. The process scales; shortcuts do not.

**Institutionalize.** Embed AI reviews into quarterly business reviews and performance conversations. Integrate proficiency into role expectations. Onboard new employees into AI-augmented workflows from day one.

## Phase 5: Optimize

Scaling is not the finish line. Optimization turns scaled deployment into compounding returns.

**Usage analytics** reveal which workflows drive the most value and where the **AI Tax**—rework from poorly integrated AI—erodes gains. **ROI tracking** translates impact into financial language the CFO cares about: Capacity Reallocation Value, cost avoidance, revenue acceleration. **Iterative improvement** deepens each 90-day cycle—early cycles build foundations, mid cycles encode knowledge, mature cycles compound advantage.

---

## Executive Sponsorship: The #1 Predictor of Scaling Success

Research consistently identifies **executive sponsorship** as the strongest predictor of whether AI pilots scale. Initiatives with strong sponsorship are **3.8x more likely to achieve objectives**. Yet only **28% of companies** report direct CEO involvement in AI governance.

Executive sponsorship is not just signing a budget approval. It is active, visible engagement: communicating why AI matters, breaking silos, resolving resource conflicts, and modeling the behavior. When [Jensen Huang tells NVIDIA employees he wants every task automated with AI](#), that is executive sponsorship. When [Meta ties AI impact to performance reviews](#), that is executive sponsorship. Without it, AI scaling dies in middle management.

---

## What Scaling Looks Like at Each Maturity Stage

The AI Maturity Model maps five stages:

**Stage 1-2 (Visibility and Adoption):** Move from scattered pilots to measured, governed deployment. Build the infrastructure and measurement that make enterprise deployment possible.

**Stage 3 (Proficiency Development):** Close the 6x productivity gap. Champions deployed. Proficiency rising. Rework rates declining.

**Stage 4 (Workflow Intelligence):** Deploy AI into workflows based on data, not opinions—mapped from how work actually flows, not how the org chart says it should.

**Stage 5 (Agentic Deployment):** AI handles primary execution paths with humans at defined checkpoints. Organizational knowledge is encoded into autonomous systems, compounding advantage over competitors running commodity tools.

---

## What the Successful Scalers Did Differently

**JPMorgan** scaled its LLM Suite to **200,000 employees**, generating **\$2 billion in business value**—built around proprietary data spanning \$10 trillion in daily transaction flow.

**Walmart** built **Wallaby** LLMs on decades of retail knowledge, scaling AI across supply chain and operations. **Shell** scaled AI monitoring to **more than 10,000 assets**, processing 20 billion sensor readings weekly. **IBM** ran “**Client Zero**” across **70+ internal workflows**, generating **\$4.5 billion in productivity gains**.

The pattern: Each organization built from the core—unique organizational knowledge—outward. The tools were the outer orbit. The knowledge was the moat. And every project had sustained executive sponsorship.

Contrast this with pilot purgatory: pilots that proved AI could work, but lacked infrastructure, proficiency development, measurement discipline, or executive backing. The technology was never the problem.

---

## The 90-Day Scaling Sprint

Wherever you are, the next 90 days follow the **Transformation Loop** from the [AI Transformation Roadmap](#).

**Days 1-14: Assess.** Audit current pilots against validation criteria. Can you name the business metric each improved? Identify which deserve to scale.

**Days 15-30: Prioritize.** Build the Phase 2 infrastructure plan. Identify champions. Select expansion targets based on workflow value and team readiness.

**Days 31-75: Execute.** Build infrastructure. Launch proficiency development. Deploy champions as mentors. Run the first expansion with full discipline.

**Days 76-90: Measure and Adapt.** Did expansion results match pilot results? What worked? What failed? Document findings and plan the next cycle.

Each cycle deepens. The first builds foundation. The second expands. The third institutionalizes. The fourth optimizes. The compounding effect is the point.

---

*Larridin provides the execution intelligence layer that makes each phase of the scaling framework measurable—from baseline adoption telemetry in Phase 1 through proficiency tracking in Phase 3 to full five-dimension impact measurement in Phase 5. If your AI pilots are working but not scaling, the gap is almost certainly in adoption depth, proficiency, governance, or measurement. That is the gap Larridin closes. [Talk to us about scaling your AI pilots](#) →*

---

# How to Build an AI Adoption Dashboard That Actually Tells You Something

You have vendor dashboards, spreadsheets, and quarterly surveys. None of them answer the question your board is asking. Here’s how to build the AI adoption dashboard your organization actually needs—from data architecture to executive visualization.

If you would like to see some of the dashboards that Larridin creates, as an example—or as an alternative to building your own—check out [the Larridin website](#). Figure 4 shows a representative dashboard from Larridin Scout.

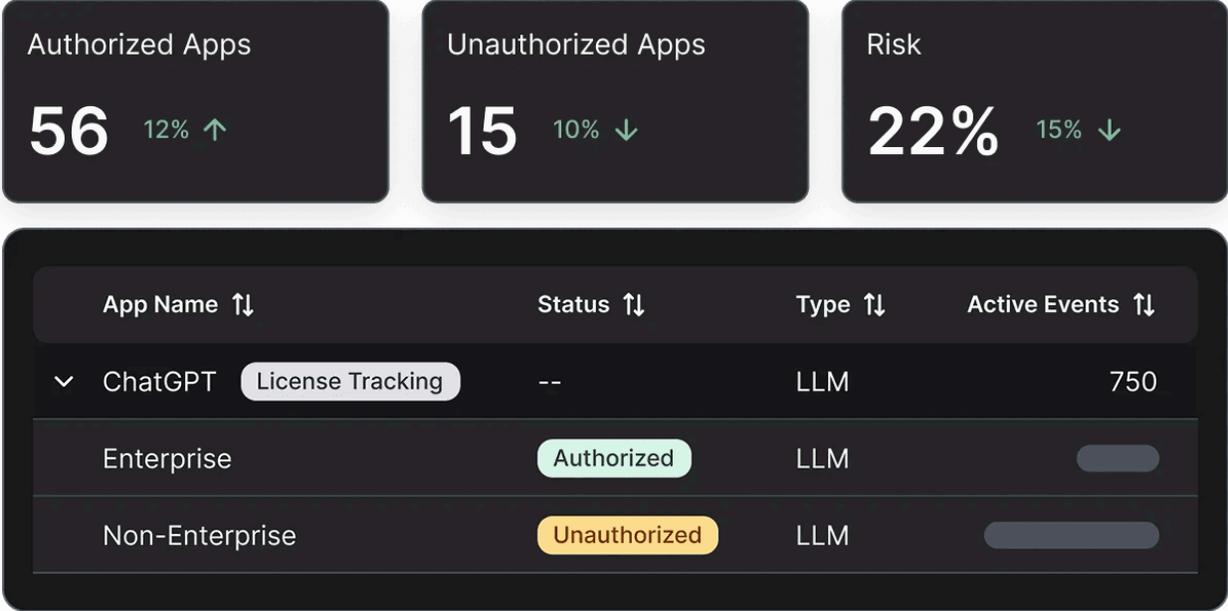


Figure 4: A representative AI dashboard from [Larridin Scout](#) (from the Larridin website)

## Spreadsheets and Surveys Don’t Scale

Every enterprise starts the same way. Someone exports Copilot usage data into a spreadsheet. Someone else pulls ChatGPT Enterprise numbers. A third person sends a survey. The results land in a slide deck that is outdated before it reaches the boardroom.

This breaks for three reasons.

First, **vendor dashboards only show their own product**. Microsoft’s Copilot Dashboard tracks adoption by group and app. Salesforce’s Agentforce analytics show agent performance. Google is rolling AI metrics into its Gemini admin dashboard. Each vendor

gives you a flattering view of their tool. None will show you what employees are doing across 15 or 30 other AI tools.

Second, **surveys measure perception, not behavior**. Employees overestimate their AI usage and underestimate rework time. Surveys alone are a foundation of sand.

Third, **spreadsheets are static snapshots**. A CSV export from two weeks ago is already wrong. You need a living system, not a quarterly data pull.

The solution is a purpose-built **AI adoption dashboard**—a single pane of glass that aggregates usage, proficiency, and impact data across every AI tool, updated continuously, segmented by organizational dimensions that make the data actionable.

---

## The Four-Layer Dashboard Architecture

The dashboard is not a flat list of metrics. It is a layered architecture where each layer answers a deeper question, mapping directly to the [AI Adoption Guide](#) framework.

### Layer 1: Tool Inventory

Before you measure anything, know what exists. The first layer is a **complete, continuously updated inventory** of every AI tool in use—sanctioned, tolerated, and unsanctioned. Display each tool's classification (AI-first, AI-augmented, agentic), modality (text, code, image, audio), governance status (approved, under review, blocked), and user count.

The IT-sanctioned list might include five tools. The actual landscape is typically 30 to 50.

### Layer 2: Usage Metrics

Layer on **quantitative usage data**: Daily, Weekly, and Monthly Active Users across every tool. Track **activation rates** (the gap between licenses purchased and licenses used), **trend lines**, and a **tool-level ranking** by engagement.

Usage metrics answer the most basic question: are people showing up? But a dashboard that stops here—and most do—is measuring deployment, not adoption.

### Layer 3: Proficiency Signals

This is where most dashboards fail. Usage tells you someone logged in. **Proficiency** tells you whether they are getting value.

Display the **adoption spectrum distribution**: Non-users, Explorers, Regular Users, Power Users, and AI-Native. Track how this shifts over time. Show **engagement trajectories** after training events. Surface top performers as **AI champions** and flag segments below the median as **intervention targets**.

This layer separates an adoption dashboard from a login counter. The [AI Transformation Guide](#) documents why this distinction—occasional user vs. AI-native—determines whether your investment generates a return or just activity.

## Layer 4: Impact Correlation

The final layer connects adoption and proficiency data to business outcomes. Do teams with higher AI adoption close deals faster? Does proficiency correlate with reduced cycle times?

This requires integration with business systems—CRM, ITSM, HRIS, and project management tools. It bridges your adoption dashboard to the [Measuring AI Impact framework](#), turning an activity report into a strategic instrument.

---

## Two Views: Executive and Operational

Build two views from the same underlying data. A single layout cannot serve the board and the AI program manager.

### The Executive View

Designed for board meetings and quarterly reviews:

- **Adoption trajectory**—Percentage of the workforce actively using AI, plotted monthly. The headline metric.
- **Department heatmap**—Adoption intensity by department, location, and hierarchy. The insight is never a company-wide average—it is where adoption is strong, weak, or absent.
- **Maturity stage indicator**—Where the organization sits on the spectrum (AI Curious through AI-Native), benchmarked against peers.
- **Governance risk summary**—Sanctioned vs. unsanctioned usage ratio, plus flagged incidents.
- **Spend efficiency**—Cost per active user. A 10,000-seat Copilot deployment with 15% weekly active usage is a spend problem, not a success story.

### The Operational View

Designed for CIOs, AI program managers, and department leads who act on the data weekly:

- **Tool-level breakdown**—Every AI tool ranked by active users, depth, and trend direction.
- **Proficiency distribution**—The adoption spectrum by team and tool.
- **Champion and laggard identification**—Champions are force multipliers. Laggards need targeted enablement, not punishment.

- **Training effectiveness**—Before-and-after engagement tied to specific events.
  - **Shadow AI detection**—Unsanctioned tools flagged in real time with user counts and data sensitivity assessments.
- 

## The Data Sources That Feed the Dashboard

Five categories of data feed a credible AI adoption dashboard.

**API logs and admin consoles.** Pull usage data from Microsoft 365 Copilot, Google Workspace, ChatGPT Enterprise, and GitHub Copilot via admin APIs. Rich vendor-specific data—but each only covers its own tools.

**SSO and SAML data.** Your identity provider (Okta, Azure AD, Ping) records every authentication event—revealing which tools employees access with enterprise credentials and which they access outside the identity layer.

**License management platforms.** Zylo, Productiv, or your SAM platform show what you pay for versus what is used. This feeds spend efficiency metrics.

**LLM gateway logs.** If you route AI traffic through a gateway, the logs provide granular data on prompt volume, model usage, and data flow patterns.

**Survey and sentiment data.** Pulse surveys capture what telemetry cannot—perceived usefulness, friction points, unmet needs.

No single source is sufficient. API logs without identity data cannot segment by department. License data without telemetry cannot distinguish paid from active seats.

---

## Key Visualizations

**Adoption curves.** Plot the S-curve of adoption over time, overall and by department. Overlay key events (tool launches, training, mandates) to show what moves the needle. The most important visualization for executive storytelling.

**Department heatmaps.** Departments on one axis, adoption depth on the other, color-coded from red (non-adoption) through yellow (exploring) to green (embedded). Leadership sees the distribution at a glance.

**Proficiency distribution histograms.** Most enterprises find a bimodal distribution—a cluster of power users and a long tail of occasional users with a gap in the middle. Closing that gap is the highest-leverage investment, as the [Copilot ROI Framework](#) details.

**Shadow AI detection views.** Unsanctioned tools with user counts, data sensitivity flags, and trend direction. This is not about catching employees—it is about understanding

demand. If 200 marketers adopted an unsanctioned design tool, that is a signal worth reading.

---

## The Unsanctioned ("Shadow") AI Problem: Your Dashboard Is Incomplete Without It

Here is the uncomfortable truth: **your AI adoption dashboard is lying to you if it only tracks sanctioned tools.**

A majority of employees use personal accounts to access AI tools IT has not approved. They paste proprietary data into public models and sign up with personal emails—not out of malice, but because approved tools do not meet their needs or procurement is too slow.

A dashboard built only on vendor APIs misses this entirely. You see Copilot adoption at 40% and conclude usage is moderate. In reality, it might be 75%—but a third is happening in tools you cannot see or secure.

Shadow AI detection is not optional. Browser-level telemetry, network monitoring, and CASB integration make it possible. The [AI Governance Guide](#) covers the governance framework in depth, but the dashboard requirement is simple: if it is not visible, it is not measured, and if it is not measured, it is not managed.

---

## Build vs. Buy

You have two paths.

**Build your own** if you have a data engineering team, BI infrastructure (Tableau, Power BI, Looker), and fewer than 10 AI tools to track. Pull from vendor APIs, pipe into your warehouse, build visualizations, maintain integrations.

**Use a purpose-built platform** when your landscape is complex (15+ tools), you need browser-level telemetry for shadow AI detection, and you want proficiency scoring out of the box. This is what Larridin does—aggregating adoption, proficiency, and impact data across every tool, team, and employee.

Most enterprises underestimate the maintenance burden of a homegrown dashboard. The initial build is straightforward. Keeping 20+ integrations current and adding new tools as they appear—that is where internal builds stall.

---

## Common Mistakes

**Tracking logins, not value.** A dashboard showing 5,000 monthly active users tells you almost nothing about whether AI is creating value. Layer proficiency and impact data on top of usage, or you are measuring the wrong thing.

**No proficiency dimension.** If every user is treated as equal—whether they ask one question a month or use AI across every workflow—your dashboard cannot guide enablement investment.

**Static snapshots instead of trends.** A dashboard showing current state without trajectory is a photograph, not a movie. Trend lines detect momentum shifts, measure training effectiveness, and catch adoption decay early.

**Measuring one tool instead of the ecosystem.** Equating AI adoption with Copilot adoption ignores the other 25 tools your employees use.

**Ignoring organizational context.** Usage data without HRIS segmentation produces numbers without meaning. “1,200 active users” is a data point. “Engineering at 92%, finance at 18%, middle management is the bottleneck” is an insight.

---

*Larridin provides the unified AI adoption dashboard described in this guide—with browser-level telemetry, shadow AI detection, HRIS integration, and reporting views for every audience from the board to the AI program manager. (See [how Larridin measures AI adoption](#) ->)*

---

## AI Adoption Benchmarks 2026: What Does “Good” Actually Look Like?

---

Gather benchmark data by industry, maturity stage, and function—so you can stop guessing where your organization stands and start measuring against reality.

### Benchmarks Measure More Than “Do You Use AI?”

The most common mistake in AI benchmarking is treating it as a binary question. “Do your employees use AI?” is a near-meaningless metric in 2026—**88% of organizations** now use AI in at least one business function ([McKinsey’s 2025 State of AI report](#)). Congratulations: you and nearly everyone else.

Real benchmarks measure three things that matter. **Depth:** not just whether employees log in, but whether AI has become habitual—daily workflows, not monthly experiments. **Breadth:** not a single chatbot, but a diverse portfolio of AI-first, AI-augmented, and

agentic tools across functions. **Proficiency**: not activity, but capability—the difference between asking ChatGPT a question once a week and integrating AI across every deliverable.

Larridin’s [Four Layers framework](#) codifies this. Layer 1 is **usage** (are people showing up). Layer 2 is **depth and engagement** (is it a habit). Layer 3 is **breadth** (how many tools and categories). Layer 4 is **segmentation** (where is adoption concentrated, and where are the gaps). Most enterprise dashboards measure only Layer 1. The organizations extracting real value measure all four.

The stakes are not abstract. [OpenAI’s 2025 State of Enterprise AI report](#) has documented a **6x productivity gap** between power users and typical employees. [McKinsey’s 2025 State of AI survey](#) found that only **6% of organizations** qualify as AI high performers—those attributing 5% or more of EBIT to AI. The gap between “using AI” and “getting value from AI” is the gap your benchmarks need to close.

---

## Industry Benchmarks: The Widening Gap

Larridin’s [AI Impact Tracker](#) monitors **567 companies** across **12 industries**, scoring each on a 1.0-5.0 maturity scale across adoption, proficiency, and impact. The industry-level data reveals an 8x gap between the leaders and laggards.

### Leaders:

Industry	% at Maturity 4.0+	Notable 5.0 Companies
Technology	74.6%	Accenture, Adobe, Alphabet, Cisco, Meta, Microsoft
Communication Services	56.5%	Netflix
Consumer Discretionary	46.1%	DoorDash
Financials	44.2%	JPMorgan Chase, BlackRock, Block, Coinbase
Healthcare	43.9%	Eli Lilly, GE HealthCare

### Laggards:

Industry	% at Maturity 4.0+
Energy	27.3%
Materials	23.1%
Real Estate	16.1%

	% at Maturity
Industry	4.0+
Utilities	9.4%

**Healthcare is the surprise performer**—43.9% above 4.0 despite heavy regulation, driven by AI’s applicability to drug discovery, diagnostic imaging, and clinical decision support. **Financial services shows the widest internal spread**: five companies at 5.0, while over 55% of the sector remains at 3.5 or below. **Utilities faces a structural challenge**—only 9.4% at 4.0, none at 4.5 or 5.0, with regulated environments and legacy infrastructure creating barriers technology selection alone cannot overcome.

If you are in a lagging industry, the competitive opportunity from early AI maturity is larger, not smaller.

## Benchmarks by Maturity Stage

The [AI Impact Tracker](#) data maps where 567 companies actually sit:

Maturity Level	% of Companies	What It Means
Below 3.0 (Nascent/Emerging)	20.7%	Pilots, limited deployment, no measurable impact
3.0-3.5 (Scaling)	38.5%	Active deployment, but proficiency and impact lag behind adoption
4.0+ (Leading)	40.9%	AI embedded in core operations with demonstrable business impact
5.0 (Transformer)	5.3%	AI-native—AI fundamentally shapes how the company competes

**The 3.0-3.5 zone is where most companies are stuck.** Larridin calls it the “scaling without depth” zone. These organizations have passed the pilot stage but haven’t built the [proficiency infrastructure](#) to convert adoption into [measurable impact](#). [Deloitte’s 2026 State of AI in the Enterprise report](#) confirms this: while **66% of organizations** report productivity gains from AI, only **20% are generating revenue** from it. The leap from 3.5 to 4.0 requires investing in proficiency development, measurement, and governance—not deploying more tools.

**The drop from 4.0 to 5.0 is steep.** Only 5.3% of companies reach the highest tier. The gap between “AI is part of how we work” and “AI is how we compete” remains the hardest to close.

---

## What the Top 5% Do Differently

McKinsey, BCG, and Larridin's tracker converge on a consistent finding: what separates the 5% achieving value at scale is not technology selection. It is execution discipline.

**They redesign workflows, not just deploy tools.** McKinsey's 2025 State of AI report found that **55% of high performers** fundamentally reworked processes when deploying AI—nearly 3x the rate of other firms. They aim for transformational change, not incremental tweaks (high performers are **3.6x more likely** to set enterprise-level transformation as the goal).

**They invest in leadership commitment.** 48% of high performers report strong senior leadership ownership of AI initiatives, versus just 16% at other companies (McKinsey, 2025). The difference is not lip service—it is structural accountability.

**They tie AI to performance.** Meta became the first major company to formally embed “AI-driven impact” into employee performance reviews in 2026. Under the policy, every employee—from engineers to marketers—is evaluated on how effectively they use AI to accelerate work and deliver results. High performers can earn bonuses of up to 200%. NVIDIA has achieved 100% AI tool adoption among software engineers. Zapier reached 97% company-wide adoption through bottom-up culture. These are not isolated examples—they are the pattern at the top.

**They measure across all four layers.** The AI Maturity Model shows that top-tier companies have complete visibility into their AI landscape, measure depth and breadth (not just logins), and segment adoption data to the team and function level.

---

## Common Benchmarking Mistakes

**Benchmarking against a single tool.** Equating AI adoption with Copilot adoption gives you a vendor-specific view, not an enterprise view. Your employees are using far more AI than any single dashboard shows.

**Using headcount instead of depth.** “5,000 employees used AI this month” tells you almost nothing. A 10,000-seat Copilot deployment with 15% weekly active usage is not a success story—it is a [spend optimization problem](#).

**Benchmarking against industry averages when you should benchmark against leaders.** If 44.2% of financial services companies are at 4.0+, the average is not your target—JPMorgan Chase at 5.0 is your target. Benchmarking against the median is a recipe for median performance.

**Ignoring proficiency distribution.** Company-wide averages hide enormous variance. If your average adoption rate is 60%, that might mean engineering is at 95% and finance

is at 20%. Without [segmentation](#), you cannot identify where enablement efforts are needed most.

**Treating benchmarking as a one-time exercise.** Adoption is dynamic. Measuring quarterly misses the trajectory. The organizations that benchmark continuously—weekly or monthly—catch stalls early and intervene before gaps widen.

---

## How to Benchmark Your Organization

Start with three questions:

1. **Can you measure adoption at all four layers**—usage, depth, breadth, and segmentation—across your entire AI tool landscape? If not, you do not yet have a foundation for benchmarking.
2. **What is your proficiency distribution?** What percentage of your workforce are non-users, explorers, regular users, power users, and AI-native? That distribution is a direct proxy for unrealized value.
3. **Where is your maturity variance?** Your engineering team may be at Stage 4 while finance is at Stage 1. A single organizational average obscures the pockets of excellence and stagnation that determine where investment will have the highest return.

The organizations that benchmark honestly—not aspirationally—are the ones that build credible board narratives, allocate resources to the right gaps, and move through the [maturity stages](#) fastest.

Larridin's [AI Maturity Self-Assessment](#) lets you score your organization across the five maturity stages and compare against the 567-company dataset from the [AI Impact Tracker](#).

---

*Larridin is the AI execution intelligence platform that gives enterprises complete visibility into AI adoption, proficiency, and impact across every tool, team, and employee. If the benchmarks in this guide revealed gaps you cannot currently measure, that is exactly the problem we solve. [Talk to us about benchmarking your organization.](#)*

---

# AI Adoption by Department: Use Cases, Metrics, and Maturity Benchmarks for Every Function

*Your company-wide AI adoption number describes no actual team in your organization. This department-by-department reference gives you the use cases, metrics, pitfalls, and maturity signals to drive adoption where it matters—function by function.*

---

## The Department Adoption Curve

AI does not spread evenly across an organization. It follows a predictable curve.

Larridin's [AI Hiring Pulse—February 2026](#) tracked 428 companies across 43,422 job postings to measure which functions are hiring for AI. The gradient is steep:

Function	Companies Hiring for AI	% of Tracked Companies
Product	81	18.9%
Customer Success	61	14.3%
Engineering & IT	54	12.6%
Data & Analytics	49	11.4%
HR & People	47	11.0%
Marketing	31	7.2%
Sales	31	7.2%
Operations	30	7.0%
Legal & Compliance	24	5.6%
Finance	20	4.7%

The gap between top and bottom is **4x**. Same organizations, same leadership, same budgets—radically different adoption intensity.

Three forces explain the ordering. **Digital-native workflows**—engineering and product work where AI slots in naturally. **Measurable outputs**—when impact is easy to quantify, investment follows. **Tool ecosystem maturity**—engineering has Copilot, Cursor, and dozens of alternatives; legal and finance have fewer proven options.

[McKinsey's 2025 State of AI report](#) confirms: 88% of organizations use AI in at least one function, but fewer than 40% have scaled beyond pilot. The curve is not a deployment problem. It is a depth problem.

---

# Department-by-Department: Use Cases, Metrics, Pitfalls, Maturity Signals

## Engineering & IT

NVIDIA reports 100% AI tool usage among software engineers. The risk here is not under-adoption—it is **tool sprawl**.

**Top 3 use cases:** Code generation and completion, automated test writing, agentic debugging. **Metrics:** Code acceptance rate, deployment frequency, defect rate in AI-assisted commits. **Pitfalls:** Tool proliferation without IT visibility; security exposure from code in unvetted models; over-reliance without review. **Maturity:** Governed tool portfolio, AI in CI/CD pipelines, measurable sprint velocity impact.

## Product

**81 companies and climbing**—the leading function in AI hiring intensity. The rising use of AI in Product signals that AI has become a **product strategy** challenge, not just a technical one.

**Top 3 use cases:** User research synthesis and feature prioritization, requirements and specification drafting, competitive intelligence aggregation. **Metrics:** Feature cycle time, research synthesis speed, AI-enabled capability coverage on roadmap. **Pitfalls:** Treating AI as a feature to ship rather than a capability to embed; AI-generated PRDs without domain validation; substituting AI synthesis for direct user contact. **Maturity:** PMs routinely using AI for research and specs, cross-functional AI feasibility assessment without full engineering dependency.

## Customer Success

**The surprise adopter**—61 companies hiring for AI in customer-facing roles, ahead of Engineering.

**Top 3 use cases:** Automated ticket triage and intelligent routing, real-time agent coaching with sentiment analysis, proactive churn prediction. **Metrics:** First-contact resolution rate, handle time, customer satisfaction (CSAT), ticket deflection rate. **Pitfalls:** Customer-facing hallucination risk; prioritizing deflection over resolution quality; brand voice inconsistency. **Maturity:** AI handling routine inquiries end-to-end with human-in-the-loop (HITL) for escalation, using a sentiment analysis feeding strategy, resolution time improving without satisfaction degradation.

## Data & Analytics

**The natural fit**—49 companies. Data teams were the earliest internal adopters because AI accelerates the work they already do.

**Top 3 use cases:** Natural language data querying and automated reporting, anomaly detection and predictive modeling, AI-assisted pipeline development. **Metrics:** Analysis

cycle time, insight-to-decision speed, model accuracy, backlog reduction. **Pitfalls:** Trusting AI insights without ground-truth validation; producing more analyses rather than better ones; neglecting data quality that AI amplifies. **Maturity:** Non-technical stakeholders querying data via natural language, routine reporting automated, analysts freed for strategic work.

## HR & People

**47 companies—outpacing both Sales and Marketing.** The intensity of AI use in HR reflects two forces: managing AI's workforce impact and deploying AI in recruiting and analytics. **Meta now evaluates AI-driven impact in performance reviews**—a policy HR must design and govern.

**Top 3 use cases:** AI-assisted screening and sourcing, skills gap analysis with personalized learning paths, attrition prediction and internal mobility analytics. **Metrics:** Time to hire, quality of hire, AI proficiency distribution across teams (see the [AI Proficiency Guide](#)). **Pitfalls:** Bias in AI-assisted hiring under intensifying regulatory scrutiny; AI in sensitive contexts without transparency; treating proficiency as IT's problem. **Maturity:** AI across the talent lifecycle, HR leading proficiency enablement, workforce planning accounting for AI-driven role redesign.

## Marketing

**31 companies.** Content production leads adoption, but strategic applications lag.

**Top 3 use cases:** Content production at scale, campaign prediction and A/B test generation, audience segmentation with personalization. **Metrics:** Content production rate and cost, campaign ROI, cost per lead, creative cycle time. **Pitfalls:** Flooding channels with undifferentiated AI content; volume over targeting; creative team resistance when AI is positioned as replacement. **Maturity:** AI embedded from brief through distribution, personalization driven by models not manual rules, team members at [proficiency Level 3+](#).

## Sales

**31 companies.** Sales teams generate massive unstructured communication. AI turns it into structured insight.

**Top 3 use cases:** Outreach drafting and call summarization, pipeline scoring and deal risk identification, competitive intelligence synthesis. **Metrics:** Deal velocity, win rate, forecast accuracy, quota attainment. **Pitfalls:** CRM data quality undermining model accuracy; seller skepticism; difficulty attributing wins to AI. **Maturity:** AI forecasts trusted by leadership, call intelligence feeding coaching, sellers adopting tools voluntarily.

## Operations

**30 companies—high potential, slow uptake.** Legacy ERP systems and cross-functional complexity create steep integration barriers.

**Top 3 use cases:** Workflow automation and process mining, demand forecasting and supply chain optimization, vendor analysis and scheduling. **Metrics:** Throughput, cost per unit, on-time delivery, process cycle time. **Pitfalls:** AI optimization layered on broken processes; underestimating legacy integration complexity; piloting in isolation from upstream workflows. **Maturity:** Process mining surfacing opportunities humans missed, demand forecasting reducing carrying costs, operations contributing to AI governance.

## Legal & Compliance

**24 companies—emerging from a low base, accelerating fast.** [Corporate legal AI adoption more than doubled in one year](#). Arista Networks' top AI-hiring function is Legal—a signal most companies have not recognized they need.

**Top 3 use cases:** Contract review and clause extraction, regulatory change monitoring, AI-assisted research and due diligence. **Metrics:** Contract review cycle time, clause accuracy, regulatory response time, outside counsel spend. **Pitfalls:** Hallucination risk where accuracy is non-negotiable; privilege concerns with cloud-based tools; treating AI governance as someone else's job. **Maturity:** Legal team governing AI usage across the organization, contract workflows incorporating AI as standard, measurable reduction in the use of outside counsel.

## Finance

**20 companies—lowest hiring intensity, yet [Gartner reports 59% of finance leaders using AI in some form](#).** The gap suggests experimentation without commitment to dedicated capability.

**Top 3 use cases:** Revenue forecasting and scenario modeling, close cycle acceleration with automated reconciliation, audit prep and expense anomaly detection. **Metrics:** Forecast accuracy, close cycle time, audit finding rate, analyst hours saved. **Pitfalls:** Zero-error tolerance creating paralysis; explainability requirements AI models struggle to meet; conservative culture treating AI experimentation as reckless. **Maturity:** AI forecasting trusted for board reporting, close cycle measurably reduced, finance positioned as AI-fluent strategic partner.

---

## Cross-Department Coordination: Avoiding the Silo Trap

The biggest risk is not under-adoption in any single function. It is **fragmented adoption across all of them**—ten departments running independent experiments with no shared governance, no common measurement, no cross-pollination.

The [AI Adoption Guide](#) calls this problem the lack of a “single pane of glass”—unified management and measurement within the company. Adopting the following three principles prevents it:

**Shared governance.** A cross-functional AI council—not just IT—owns tool sanctioning and data policy. Without shared governance, every department builds its own (usually inadequate) shadow AI ecosystem.

**Common measurement.** The Four Layers framework—Usage, Depth, Breadth, Segmentation—provides consistent metrics across functions. When engineering measures “daily active users” and marketing measures “content pieces generated,” you cannot compare or allocate.

**Cross-functional learning.** The department that solves a hard problem—Customer Success managing hallucination risk, say—has lessons that transfer directly to Marketing, Sales, and HR. Build the mechanisms: case studies, working groups, shared prompt libraries by function.

The companies at the top of the AI Hiring Pulse—Moody’s, NVIDIA, Cisco, Kraft Heinz—treat AI as an **organizational capability**, not a collection of departmental initiatives. That is the gap.

---

*Larridin maps AI adoption, proficiency, and impact at the department, team, and individual level—across every AI tool in your environment. Stop reporting company-wide averages. [See your department-level adoption data](#) →*

---

## How to Drive AI Adoption in Non-Technical Teams: A Department-by-Department Playbook

---

The next wave of enterprise AI value will not come from engineering. It will come from Sales, Marketing, HR, Finance, Legal, and Operations—where AI adoption is lowest and unrealized returns are highest.

### The Adoption Gap Is Real—and Widening

Most enterprise AI investment flows to engineering, product, and IT. Most enterprise value does not live there. Sales, Marketing, HR, Finance, Legal, and Operations represent the majority of your workforce and cost structure—and they have the lowest AI adoption rates in the organization.

The data is stark. [Worklytics benchmarks](#) show engineering teams in mature organizations operating at **80-90% meaningful AI adoption**. Business functions sit at **20-40%**. [McKinsey’s 2025 State of AI report](#) confirms the pattern: adoption is highest in IT and product development, significantly lower in HR, legal, and finance. [EY’s 2025 Work Reimagined Survey](#) found that while 88% of employees report using AI daily, **only**

**5% use it in advanced ways**—and that 5% clusters overwhelmingly in technical functions.

This is a value gap. [McKinsey's economic potential of generative AI research](#) estimates that sales and marketing alone account for **28% of the total economic value** from generative AI—more than any other function. The departments with the most to gain are adopting the slowest. The [AI Adoption Maturity Spectrum](#) shows most non-technical teams stuck at Stage 1 or Stage 2. This guide is a playbook for closing that gap.

---

## Why Business Teams Resist—and Why They Are Different From Tech

The adoption gap is not about intelligence or capability. It is structural, and the barriers are fundamentally different from those in engineering.

**Fear of replacement, not fear of change.** Engineers see AI as a power tool. Non-technical employees read the headlines differently. [BCG's AI at Work 2025 report](#) found that employees at organizations undergoing AI-driven redesign are **significantly more worried about job security**—46% versus 34% at less-advanced companies. And 37% worry that overreliance on AI will erode their skills entirely. This fear is not irrational. It has to be addressed directly, not dismissed.

**Unclear value proposition.** Engineers know exactly what to do with an AI coding assistant. A VP of Sales does not have an equally obvious starting point. When the use case is not self-evident, adoption stalls at curiosity. The [AI Proficiency Guide](#) shows non-technical teams disproportionately stuck at **Level 1—Search Replacer**—because nobody has shown them what Level 2 or 3 looks like in their context.

**Training designed for the wrong audience.** A prompt engineering workshop centered on code generation is useless for an HR director streamlining recruiting workflows. [BCG reports](#) that less than 25% of employee learning time happens during work hours—AI skill-building gets bumped to personal time. Without protected, role-relevant training, non-technical employees default to basic usage and never progress.

**No measurement, no accountability.** If you do not measure adoption at the department level, you cannot manage it. A 60% company-wide adoption rate can mean engineering is at 95% and Legal is at 15%. Without segmented data, the gap is invisible.

---

## The Opportunity Map: Department by Department

Every business function has high-value AI use cases. The gap is not opportunity—it is awareness and enablement. Here is where the value sits.

**HR: Recruiting, Performance Management, Workforce Analytics.** Resume screening is the entry point, but deeper value is in onboarding personalization, skills gap analysis, attrition prediction, and performance review calibration. HR is uniquely positioned to lead organizational AI transformation—not just support it. The [AI Transformation Guide](#) makes this case in detail.

**Marketing: Content, Analytics, Personalization.** Content production at scale is the obvious starting point. Campaign optimization, audience segmentation, and competitive messaging analysis are where analytical depth emerges. Content gets marketing through the door. Analytics keeps them there.

**Sales: Pipeline Intelligence, Call Analysis, Forecasting.** Proposal generation, pipeline scoring, deal risk identification, call summarization, and win/loss pattern recognition. Sales teams produce enormous unstructured communication. AI turns that into structured insight.

**Finance: Forecasting, Audit, Reporting.** Scenario analysis, close cycle acceleration, audit preparation, anomaly detection, and board reporting. Finance resists AI because zero tolerance for error is cultural—but AI excels at exactly the pattern-recognition tasks that consume analyst hours.

**Legal: Contract Review, Compliance Monitoring.** Clause extraction, regulatory change tracking, due diligence, and policy drafting. Legal handles some of the most text-intensive work in any organization—a natural fit for language-based AI.

**Operations: Process Optimization, Supply Chain.** Workflow automation, demand forecasting, supply chain monitoring, and SOP generation. High potential, slow uptake due to legacy ERP systems and fragmented tooling.

The section on AI adoption by department provides detailed profiles, tools, metrics, and blockers for each function.

---

## The Champion Model: Your Most Scalable Lever

AI adoption in non-technical teams fails when it is driven by IT mandates. The champion needs to be someone **inside the department**—the marketing manager who cut content production time in half, the recruiter who built an AI-assisted screening workflow, the analyst who automated report narratives.

These champions have credibility that an IT-imposed rollout never will. They speak the department's language, understand the workflows, and signal to peers that AI is a competitive advantage—not a threat. [Writer's 2025 enterprise AI adoption report](#) found that **77% of employees using AI already self-identify as champions or see the potential to become one**. The talent is there. It needs to be identified and empowered.

How to build the network: **Identify through data, not self-nomination**—use behavioral adoption data to find Level 3-4 users in each department. **Give them platforms**—let champions run department-specific workshops and share workflows. **Make it visible**—recognize champions in performance reviews and internal communications. **Scale horizontally**—one champion per department is a start, but the goal is a network that reaches every team.

Zapier achieved **97% AI adoption across all departments** by pairing top-down urgency with bottom-up building. They ran a full-company hackathon, got every team building with AI, published clear guidelines, and appointed champions who helped others get started. CEO Wade Foster’s key insight: the best unlocks come from the edge, where the work is being done, not from top-down mandates.

---

## Case Study: Meta Ties AI Adoption to Performance Reviews

In early 2026, **Meta became the first major company to formally tie employee performance reviews to AI usage**—across all roles, not just engineering. Under the new policy, **“AI-driven impact” is a core performance expectation** for every employee, from engineers to marketers. Managers evaluate how effectively employees leverage AI to accelerate work and deliver results.

Meta’s Head of People, Janelle Gale, stated that the company wants to recognize people helping it move toward an AI-native future. To support the mandate, Meta built internal tooling including a gamified system called **Level Up** that rewards employees with badges as they hit AI milestones—encouraging experimentation before formal evaluation begins.

You do not have to go as far as Meta on day one. But you need to connect AI adoption to something that matters to the individual—performance reviews, promotion criteria, team goals, visible recognition. Without aligned incentives, the gap between technical and non-technical teams persists.

---

## Training That Works for Non-Technical Users

Generic “prompt engineering bootcamps” do not work for business teams. Here is what does.

**Start with one killer use case per department.** Find a task that is painful, repetitive, and immediately improved by AI. For Sales, it might be first drafts of proposals. For Legal, contract clause extraction. For Finance, variance commentary on monthly reports. The goal is to find key moments where a non-technical employee says, “I cannot go back to doing this the old way.”

**Build role-specific training, not tool-specific training.** Do not train people on ChatGPT. Train the marketing team on competitive analysis workflows. Train finance on scenario modeling. Train legal on contract review. The tool is the vehicle, not the focus.

**Make learning continuous, not one-time.** Monthly tool briefings. Curated use-case libraries updated regularly. Peer learning networks where employees share discoveries. The AI landscape changes too fast for a single training event to hold value beyond a few weeks.

**Protect time for experimentation.** Most AI learning happens outside work hours—which means it barely happens at all. Allocate structured experimentation time: AI happy hours, hackathons, and dedicated sessions that create the psychological safety non-technical teams need to progress beyond Level 1. The [AI Proficiency Guide](#) details how to measure whether training is actually moving people up the proficiency spectrum.

---

## Measuring Adoption in Business Teams—Different Signals

The metrics that work for engineering—code acceptance rates, deployment frequency, AI-assisted commits—are meaningless in a marketing or HR context. Business team adoption requires different signals.

**Track department-level adoption rates independently.** Stop reporting company-wide averages. Segment by department, team, and use case. The section on AI adoption by department provides the framework.

**Measure proficiency distribution, not just usage.** A department where 80% of people use AI at Level 1 is fundamentally different from one where 40% use it at Level 3. Use the [AI Proficiency Guide](#) to understand not just who is using AI, but how well.

**Look for workflow integration, not login counts.** The real signal is whether AI is embedded in how a team operates: are proposals drafted with AI by default? Are contracts reviewed through AI-assisted tools as standard practice? Workflow integration means adoption has moved from experimentation to habit.

**Surface wins invisible in company-wide data.** A finance team reducing close cycle time by two days through AI-assisted reconciliation is significant. But wins like this one disappear in a company-wide metric. Department-level data makes these wins visible—and visibility drives momentum.

---

*Larridin gives enterprises complete visibility into AI adoption across every department—not just engineering. Our platform measures adoption at the team level, identifies champions, surfaces the proficiency gap between functions, and shows exactly where enablement will have the highest return. [See how Larridin drives adoption across your entire organization.](#)*

---