melt

MELT RESEARCH REPORT · OCTOBER 2026

Everyone is building AI. Who is rebuilding the business?

Companies are investing in AI and hiring leaders to drive its adoption. But getting AI to work across the business requires more than technology. It means making knowledge accessible, demonstrating business value, and establishing clear ownership of what gets built.



Contents

01 The knowledge gap 8
AI can read your documents. It can't access the experience behind your decisions.

02 The value gap 11
Easy to approve. Hard to prove.

03 From prototype to production 14
A working agent can be built in a single night. Making it production-ready is where the real work begins.

04 The role 16
Everyone agrees the role is needed. No one agrees what it is.

05 Who owns AI 19
Who decides, who builds, and who answers for the results?

WHY WE WROTE THIS REPORT

Melt works with operations, supply chain and manufacturing teams. We focus on the decisions that keep a plan on track when reality changes: a supplier moves a date, a customer changes an order, a line runs slow.

To build for those teams, we spent a lot of time learning how their companies are adopting AI. Along the way we kept meeting the same work in other industries: leaders hired to drive AI transformation, facing the same open questions. So we widened the lens, and this report shares what we found about AI transformation as a whole.

What we found goes beyond any one industry. Companies may operate differently, but the challenges of turning AI ambition into business results are remarkably similar: making knowledge accessible, demonstrating value, and working out who is responsible for what gets built. These are organizational questions as much as technical ones, and they are still being worked out.

Michal Gil and Omer Schneider
Melt

SUMMARY

AI can now build in a night what organizations take months to absorb. Prototypes are cheap, but making them production-ready takes far longer. Much of the know-how AI agents need is held by employees or in systems the agents cannot reach. AI leaders are asked to weigh proposals from every part of the business, and no one can be an expert in all of them; proving the results afterward is harder still. And the role created to fix all this is itself still undefined.

Key findings

1

CHAPTER 1

The knowledge gap: access to information is not the same as access to context.

Practitioners described critical business knowledge as scattered across undocumented processes, employee experience and private conversations. Yet none of the job postings asks the AI transformation role to make that knowledge more accessible. Changing how people share and maintain knowledge was consistently described as harder than building the tool.

2

CHAPTERS 2 AND 3

The value gap: AI leaders must choose initiatives they cannot fully judge, and then prove their value.

The people choosing initiatives cannot be experts in every domain, so they rely on estimates from the teams proposing them. Several practitioners we interviewed track usage and cost, but lack an objective way to measure the business value leadership asks for. Most raised demonstrating ROI as an unresolved challenge.

3

CHAPTERS 4 AND 5

The ownership gap: companies are still working out who should own AI, and no standard model has yet emerged.

Organizations are experimenting with departmental champion networks, business ownership of agents and new governance routines. Even the role meant to lead this work is undefined: employers title, level and scope it in very different ways.

INTRODUCTION

Building AI is getting easier. Organizing a business around it is not.

Companies are hiring people to lead AI transformation, establish new operating models and help the business adopt AI. But getting AI into an organization is not the same as changing how the organization works.

We conducted more than 100 interviews with practitioners involved in AI transformation, including AI leaders, CIOs, operations leaders and independent consultants, and examined over 100 roles companies are hiring for. Three challenges kept emerging: making organizational knowledge accessible to AI, choosing initiatives and proving their business value, and deciding who owns what gets built.

These responsibilities reach across the business, but the authority, expertise and accountability needed to deliver them are rarely concentrated in one place. The result is a role whose scope is expanding faster than its definition. The AI transformation leader is a window into the larger organizational challenge: companies are creating a role to coordinate a transformation that no single role, team or function can deliver alone.

This report examines the three gaps behind that tension, how organizations are addressing them, and what they reveal about the evolving role of the AI transformation leader.



FRAMEWORK

The AI journey

To place the organizations we spoke with, Melt uses its working model of the AI journey (Exhibit 1), informed by current research and extending into anticipated stages. In the model, each stage addresses the friction of the one before it and introduces its own. Many organizations are confronting the value bottleneck, the third stage in the model: initiatives multiply faster than their impact can be demonstrated. None of the organizations we interviewed had yet reached the fourth stage. The three gaps this report examines are most visible at the value bottleneck.

What moving to the next stage looks like. An industrial IoT company illustrates the shift from the second stage to the third. Employees across the company had been building their own AI tools, and some of it paid off: an employee-built technician scheduling tool is now integrated into the company's core systems. But the company also discovered that seven teams had independently built the same AI skill for creating company presentations. Rather than restrict bottom-up building, it is adding structure around it: a shared repository from which employees can pull existing skills instead of recreating them, and common standards for how skills are built. Another company tackles the same problem with a recurring builders' forum, where people building with AI in any department share and coordinate their work.

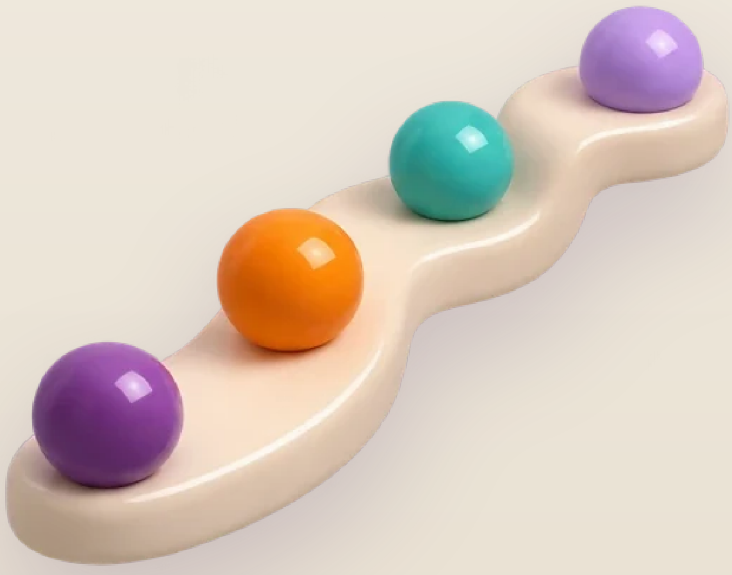


EXHIBIT 1

Many organizations are confronting the value bottleneck: initiatives multiply faster than their impact can be demonstrated.

The AI journey, Melt framework

STAGE		WHAT HAPPENS	FRICTION	EXAMPLES FROM THIS RESEARCH
1	**R&D Playground**	Engineering experiments with AI in isolation.	Disconnected from the business.	One AI consultancy reports that most of its clients are at this stage, adopting developer productivity tools.
2	**Shadow Builds**	Every business unit builds its own.	Nobody knows what exists.	Seven teams at an industrial IoT company built the same presentation skill. At a technology company, the AI transformation lead found that a senior employee had invested significant time building an internal social app that addressed no business need and delivered little value.
3	**The Value Bottleneck** WHERE COMPANIES ARE STUCK	Organizations introduce AI offices, review gates or other structures to impose order.	Too many initiatives, no ROI clarity.	Several organizations had set up an AI office or a review gate to bring order to their initiatives. Even with that structure in place, leaders struggled to say which initiatives were paying off. Lacking an objective way to measure value, they often have no option but to rely on each team's own estimate of its impact.
4	**Agents at Work** ANTICIPATED	Agents are deployed widely and continually enhanced.	Expected: agents drift, and knowledge gets trapped inside individual agents.	Not yet reached by the organizations in this research.
5+	**The Adaptive Business** ANTICIPATED	Thousands of agents and humans, running as one governed operating model.	Not yet observed.	Not yet reached by the organizations in this research.

Source: Melt practitioner interviews, 2026. Stages assigned by Melt, not self-reported.

AI can read your documents. It can't access the experience behind your decisions.

Connecting AI to company data is increasingly straightforward. The harder problem is giving it the knowledge people use to make decisions: the exceptions, trade-offs and unwritten rules that rarely make it into company systems.

FROM THE FIELD

Semiconductor-equipment manufacturer

Planners constantly reconcile changing demand, supply constraints and customer commitments. The sales forecast changes every two weeks, triggering thousands of changes to orders, dates and supplier coordination.

When materials run short, they decide which customers get priority, which orders can wait and where production should be reshuffled.

THE KNOWLEDGE GAP

The ERP records what was decided, not why.

Those decisions rely on experience, exceptions and trade-offs that rarely make it into company systems.

01 · THE KNOWLEDGE GAP

Access to information is not the same as access to context. An agent can read every document a company holds and still lack what a ten-year employee knows.

This was among the most consistent themes in our interviews. Much of the knowledge needed to run the business is undocumented, held in employees' heads, or kept in private messaging channels that company-wide AI tools cannot access. Even with full access to documents and databases, an agent lacks the context of a long-tenured employee.

One technology company put a number on it. Its transformation lead estimated that only a single-digit share of internal Slack communication took place in public channels. The rest happened in private channels and direct messages, often inaccessible to colleagues outside those conversations and unavailable to company-wide AI tools.



The gap is behavioral as much as technical. Making context available to AI depends on how people work: whether discussions happen in shared channels rather than private messages, and whether processes are written down. One company is encouraging employees to move discussions into public channels and to keep a single source of truth for exactly this reason. Another made documenting its processes the first task of its new AI lead, ahead of any agent deployment.

Yet job descriptions do not ask the AI role to solve it. None of the postings with a usable description asks the role to capture employee knowledge (see Exhibit 4 in Chapter 4). Where they refer to knowledge at all, they mean knowledge bases and search tools, or documenting what the AI program itself has learned.

THE TAKEAWAY

The knowledge AI needs most is held by people and in private channels, and none of the AI transformation job postings we analyzed asks the role to capture it. That responsibility may sit with knowledge management, IT or operations. Leaders should make sure it sits somewhere.

IMPLICATIONS FOR LEADERS

01 Define where organizational knowledge should live as part of any AI readiness assessment.

02 Assign clear ownership for capturing it, whether in the AI team, knowledge management, IT or operations.

03 Encourage working practices that leave important context behind in shared, appropriately accessible channels and documented processes.

Easy to approve. Hard to prove.

No one can be an expert in every area AI touches, yet the AI leader is asked to judge proposals from all of them, and then to prove what the chosen initiatives delivered.

No one person can judge every proposal. AI initiatives come from every part of the business, from finance and operations to HR and sales, so AI leaders depend on the teams proposing the work to explain what it will deliver.

In practice, the estimate of an initiative's value usually comes from the people who want it approved. But even with a structured process, expected returns are hard to verify, and practitioners described decisions sometimes going to whoever argues hardest.

Even with expertise, the comparison is difficult. Initiatives in different functions promise different kinds of benefit, such as hours saved in one, faster delivery in another and new revenue in a third, and they are hard to weigh against each other. To make these choices more consistent, the organizations we interviewed have developed structured methods for evaluating AI opportunities (Exhibit 2).

EXHIBIT 2

Organizations have developed structured methods for choosing which AI initiatives to pursue.

Four approaches to selecting AI initiatives, as described by the organizations using them

01

Opportunity heat map

LARGE MULTINATIONAL COMPANY

Every opportunity is scored on time to value, expected ROI, data readiness and delivery readiness. The scores separate quick wins from larger bets and guide where resources go.

02

Three-outcome test

SPORTS TECHNOLOGY COMPANY

A department's proposal is approved only if it promises at least one of three outcomes compared with current effort: lower cost, new revenue or faster delivery.

03

Pass/fail impact gate

TECHNOLOGY COMPANY

A prototype moves to production only after it has shown impact in a small local experiment.

04

Impact-versus-effort ranking

INDEPENDENT STRATEGY CONSULTANT

Initiatives are ranked by effort and speed against their strategic importance, with value expressed in hours saved, to produce a sequenced roadmap for the board.

Approaches are listed, not ranked; the evidence does not indicate that any one is more effective.
Source: Melt practitioner interviews, 2026

The unresolved challenge comes after launch. Measurement was the most frequently cited unresolved challenge in our interviews, raised by organizations at every stage of maturity. The difficulty is translation.

AI'S EFFECTS ARE USUALLY RECORDED AS

hours saved, tasks automated or tools adopted.

LEADERSHIP MEASURES THE BUSINESS THROUGH

margin, revenue, cost and on-time delivery.

Many organizations lack a reliable way to connect the first to the second. The AI leader of a large multinational company had built a scoring model for every opportunity, yet identified demonstrating ROI to the CEO as the most important unmet need.

Some practitioners stop at the first step and report hours saved, a figure that is easier to defend but does not answer leadership's question.

THE TAKEAWAY

The AI leader is asked to weigh proposals from every part of the business, relying on estimates from the teams proposing them, and then to prove results in terms the business has never defined for AI.

IMPLICATIONS FOR LEADERS

01 Before approving an initiative, ask its business owner, not the AI team, to name the business KPI it should move.

02 **Have a method for confirming results after launch,** not only for selecting initiatives.

03 Make the AI leader responsible for comparable proposals and consistent measurement, not for expertise in every domain.

A working agent can be built in a single night. Making it production-ready is where the real work begins.

Scaling a prototype is the hard part. Practitioners emphasize that building is the shorter part of the process. The work that follows includes compliance review, API approvals, integration with existing systems, security, and the reliability needed for use beyond the original builder.

One AI transformation lead put numbers on it: an agent can reach about 80 percent of its functionality in a single night, but making it production-grade can take weeks of work from one or two engineers.

Introducing AI one step at a time reveals which parts of a process should change

Getting AI into production is not simply a technical milestone. It is where organizations begin to discover what needs to change in the way work gets done. A fintech company did not redesign its customer onboarding before automating it. It embedded agents in individual steps and adjusted the workflow as it learned (Exhibit 3).

EXHIBIT 3

At one fintech, automating a single onboarding step showed that the step itself was unnecessary.

How customer data collection changed during onboarding, as described by the company

1 — 2 — 3

BEFORE

A manual, four-step process

1. The team sends the customer a spreadsheet of required data fields
2. The customer fills it in
3. The team analyzes the data
4. The team builds a data dictionary by hand

AI ADDED TO ONE STEP

An agent automates parts of the work

Automating it shows the team that the customer should not need to fill in the spreadsheet at all.

THE PROCESS CHANGES

The spreadsheet step is dropped

The workflow moves toward collecting the data automatically. Several other onboarding steps also prove unnecessary and are removed rather than automated.

Source: Melt practitioner interviews, 2026

The biggest gain came from what automation exposed. Once data collection was automated, customers no longer needed to fill in the spreadsheet at all, and several other onboarding steps proved unnecessary and were dropped rather than automated. The redesign followed from evidence rather than from an up-front plan.

THE TAKEAWAY

Prototypes are quick to build. Making them production-ready takes the real work. Automating one step at a time can also reveal which process steps can be simplified or eliminated.

IMPLICATIONS FOR LEADERS

01 Budget and staff for what comes after the prototype: validation, integration, approval and ongoing ownership.

02 Decide early whether an initiative warrants production or should be bought.

03 Introduce AI one step at a time, and use what each step reveals to decide whether that step should exist at all.

Everyone agrees the role is needed. No one agrees what it is.

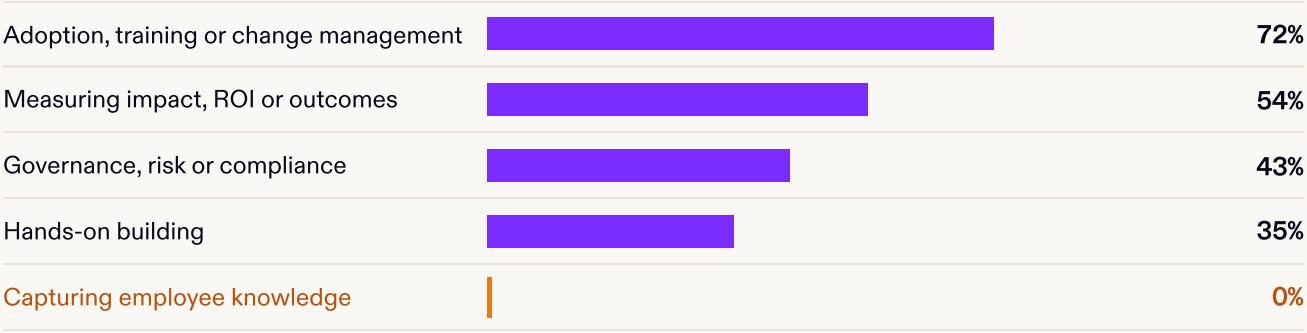
The role meant to close these gaps is itself still being defined. Employers broadly agree that AI transformation requires a dedicated owner, but they have not converged on the role's title, seniority, compensation or scope.

The job is broadly the same. Read side by side, the postings describe much the same work: driving adoption, measuring impact, establishing governance and building AI tools (Exhibit 4).

EXHIBIT 4

Most postings call for adoption and measurement; none refer to capturing employee knowledge.

Postings that mention each theme, % of postings with a usable description (n = 94)



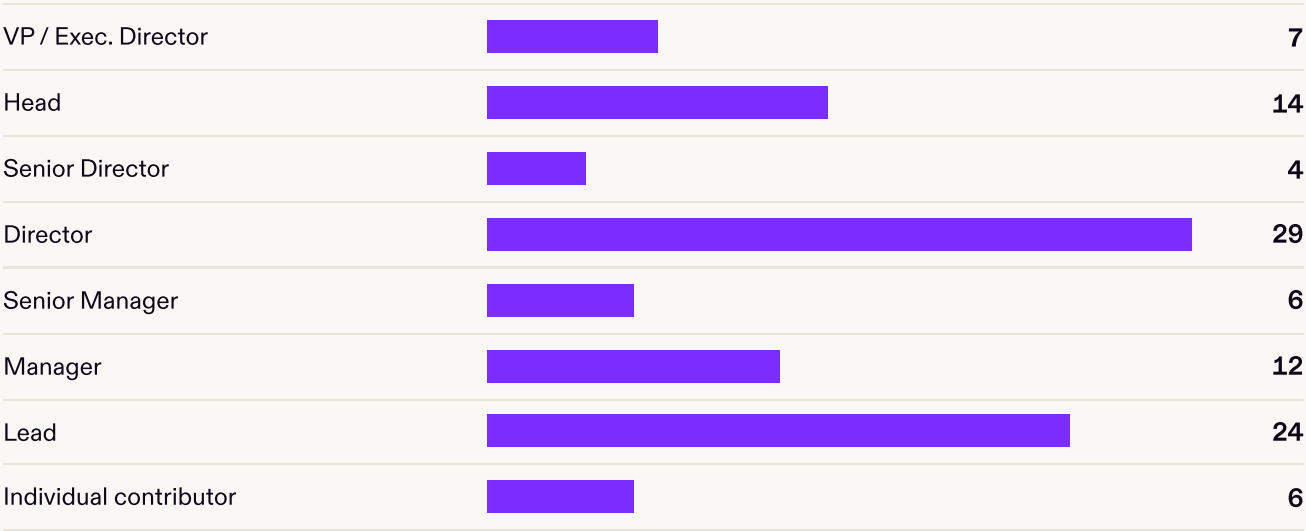
Source: Melt analysis of public job postings, July-October 2026

Yet the same work is advertised at eight levels of seniority, from individual contributor to vice president (Exhibit 5).

EXHIBIT 5

Broadly similar AI transformation roles are advertised at eight levels of seniority, from individual contributor to vice president.

AI transformation roles by seniority, number of postings (n = 102)



Includes only roles whose title or description centers on AI transformation.

Source: Melt analysis of public job postings, July-October 2026

Pay varies just as widely. The same title can pay twice as much at one company as at another, and a Lead at one company may earn more than a Director at the next (Exhibit 6, in the appendix). Location, company size and scope explain part of this, but the spread is wide enough that the title alone says little about the job.

$77K-$351K

The range of advertised base salaries across the 39 postings that disclosed pay.

Salary ranges overlap substantially across seniority levels, suggesting that titles alone offer little indication of how companies value AI transformation roles.

39 of 102 postings disclosed pay. Full list: Exhibit 6, page 27. Source: Melt analysis of public job postings, July-October 2026

The scope is often more than one person can carry. A recurring theme in our interviews was that companies expect a single hire to cover strategy, discovery, architecture, implementation and change management, a scope that would typically require a team.

Whether the role will last is an open question. One practitioner expects the dedicated role to disappear, or to change fundamentally, once a company becomes AI-native. A plausible outcome is that the program-level role diminishes, while the work of deciding what to automate and overseeing agents becomes part of how each function operates.

WHAT THE EVIDENCE SUGGESTS

Companies have decided this work needs an owner before deciding what the work is. Each is designing the role independently, which helps explain why titles, levels and pay vary so widely.

IMPLICATIONS FOR LEADERS

01 **Define the job by what it must deliver, not by its title.** Decide which outcomes the role owns before deciding what to call it or what level it sits at.

02 **Decide its reach.** Will it cover the whole company or a single department? The job is very different in each case.

03 **Plan for a team, not one person.** Strategy, architecture, implementation and change management are usually more than one hire can cover.

04 **Don't rely on titles or salaries from other companies.** The same title describes very different jobs.

Who decides, who builds, and who answers for the results?

Several organizations no longer try to run AI from one central team. Each department owns the agents it uses, and a small central team sets standards and keeps everyone coordinated. It is an emerging pattern, not yet a proven model.

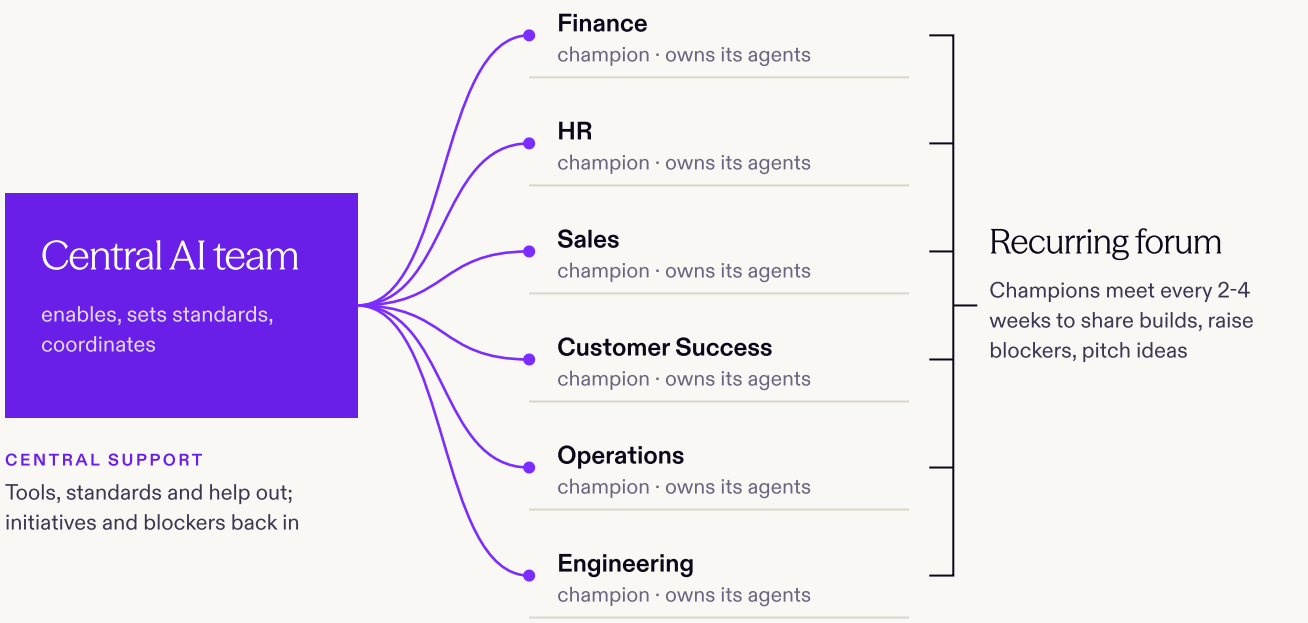
The structure described most often has three components: a small central AI team that enables rather than builds, a designated champion in each department, and a recurring forum.

At a sports technology company, AI champions spend around 10 percent of their time on AI initiatives. They remain within their own departments rather than reporting to the AI lead, and come together every two weeks. A fintech company holds weekly one-to-one meetings with its champions and a monthly group session. A large multinational company operates an AI office with representatives from departments such as HR, Legal and Finance.

EXHIBIT 7

An emerging operating model pairs a small central team with departmental owners and a recurring forum.

Recurring elements of the operating model described by interviewees



A synthesis of practices observed at several organizations; not a universal blueprint. Meeting cadence and reporting lines vary.

Source: Melt practitioner interviews and analysis of public job postings, 2026

05 · WHO OWNS AI

“Every agent has a manager, a cost, a performance review, and a forward plan.”

FROM AN AI TRANSFORMATION JOB POSTING

The model is also appearing in job requirements. Two employers ask the role to train and mentor departmental AI champions. A third employer's lead manages "dotted line champions" alongside direct reports. A networking hardware company expects its marketing AI lead to establish and chair an internal AI council. Another goes further, asking its AI transformation lead to ensure that "every agent has a manager, a cost, a performance review, and a forward plan." This indicates that the practice is spreading into role definitions; it does not establish that it outperforms alternative models.

Most of this ownership sits inside the company. Seventy-two percent of the AI transformation roles we analyzed drive adoption within the employer's own organization. The remaining 28 percent deliver AI transformation to clients, at consultancies, cloud providers and software companies that are repositioning customer-success teams as transformation partners (Exhibit 8).

EXHIBIT 8

Nearly three-quarters of AI transformation roles focus on the employer's own adoption.

Postings by orientation of the role, % of classifiable postings (n = 100)



Internal: drives AI adoption within the employer's own organization.

Client-facing: delivers AI transformation for the employer's customers.

This 72% is a different measure from Exhibit 4, where 72% of postings with a usable description (n = 94) mention adoption, training or change management.

Source: Melt analysis of public job postings, July-October 2026

THE TAKEAWAY

A central team can enable AI adoption, but it cannot judge whether an HR or Finance agent is producing correct output. In the organizations we studied, ownership is moving to those who can, and the central team's role is shifting toward standards, coordination and oversight.

IMPLICATIONS FOR LEADERS

01 **Name an owner in the business before an agent is deployed,** rather than launching without one.

02 **Set up a builders' forum and a shared library early,** before teams build the same tools in parallel.

03 **Budget for maintenance and review agents regularly.** Launch is not the finish line.

THE NEXT QUESTION

How should this operating model work?

This report has examined the shift of AI ownership toward the departments that use the agents and understand the processes behind them. But distributed ownership creates its own set of challenges.

Not every department has the technical capability to build and maintain agents. Business teams validating their own agents raises questions about independent testing. And when agents operate across multiple departments, ownership becomes harder to define.

Who owns an agent that serves several functions? Who is responsible for testing it independently? How can organizations distribute ownership without creating fragmented, incompatible AI systems?

These are questions of operating-model design. They deserve a closer examination of their own, and will be the subject of a separate Melt research report.



The organizations in this study did not lack AI ambition or experimentation. What they lacked, to varying degrees, was the organization around it: clear ownership, accessible knowledge, credible measures and established routines. The next phase of AI transformation is not just about bringing AI into the business. It is about reorganizing the business so AI can operate within it: with access to the right knowledge, a clear measure of value, and someone accountable for the outcome.

ABOUT THE RESEARCH

This report draws on two sources. An analysis of over 100 AI transformation job postings shows how employers describe, level and compensate the role. Interviews with AI transformation leaders and independent consultants show what practitioners do and where they encounter obstacles.

AUTHORS · MICHAL GIL AND OMER SCHNEIDER, MELT RESEARCH

102

AI transformation job postings

Includes only roles whose title or description centers on AI transformation. Source: Melt analysis of public job postings, July-October 2026.

100+

Practitioner interviews

Interviews with AI transformation leaders and independent consultants. Source: Melt practitioner interviews, July-October 2026.

The interviews span over 100 organizations across Israel, the US and Europe. About 40% are in manufacturing, supply chain and other operations-heavy businesses, including food and agricultural producers, electronics and semiconductor makers, and retailers; the rest come from software, telecom, financial and professional services. Interviewees ranged from heads of AI transformation and CIOs to operations leaders and independent consultants.

ABOUT THE RESEARCH

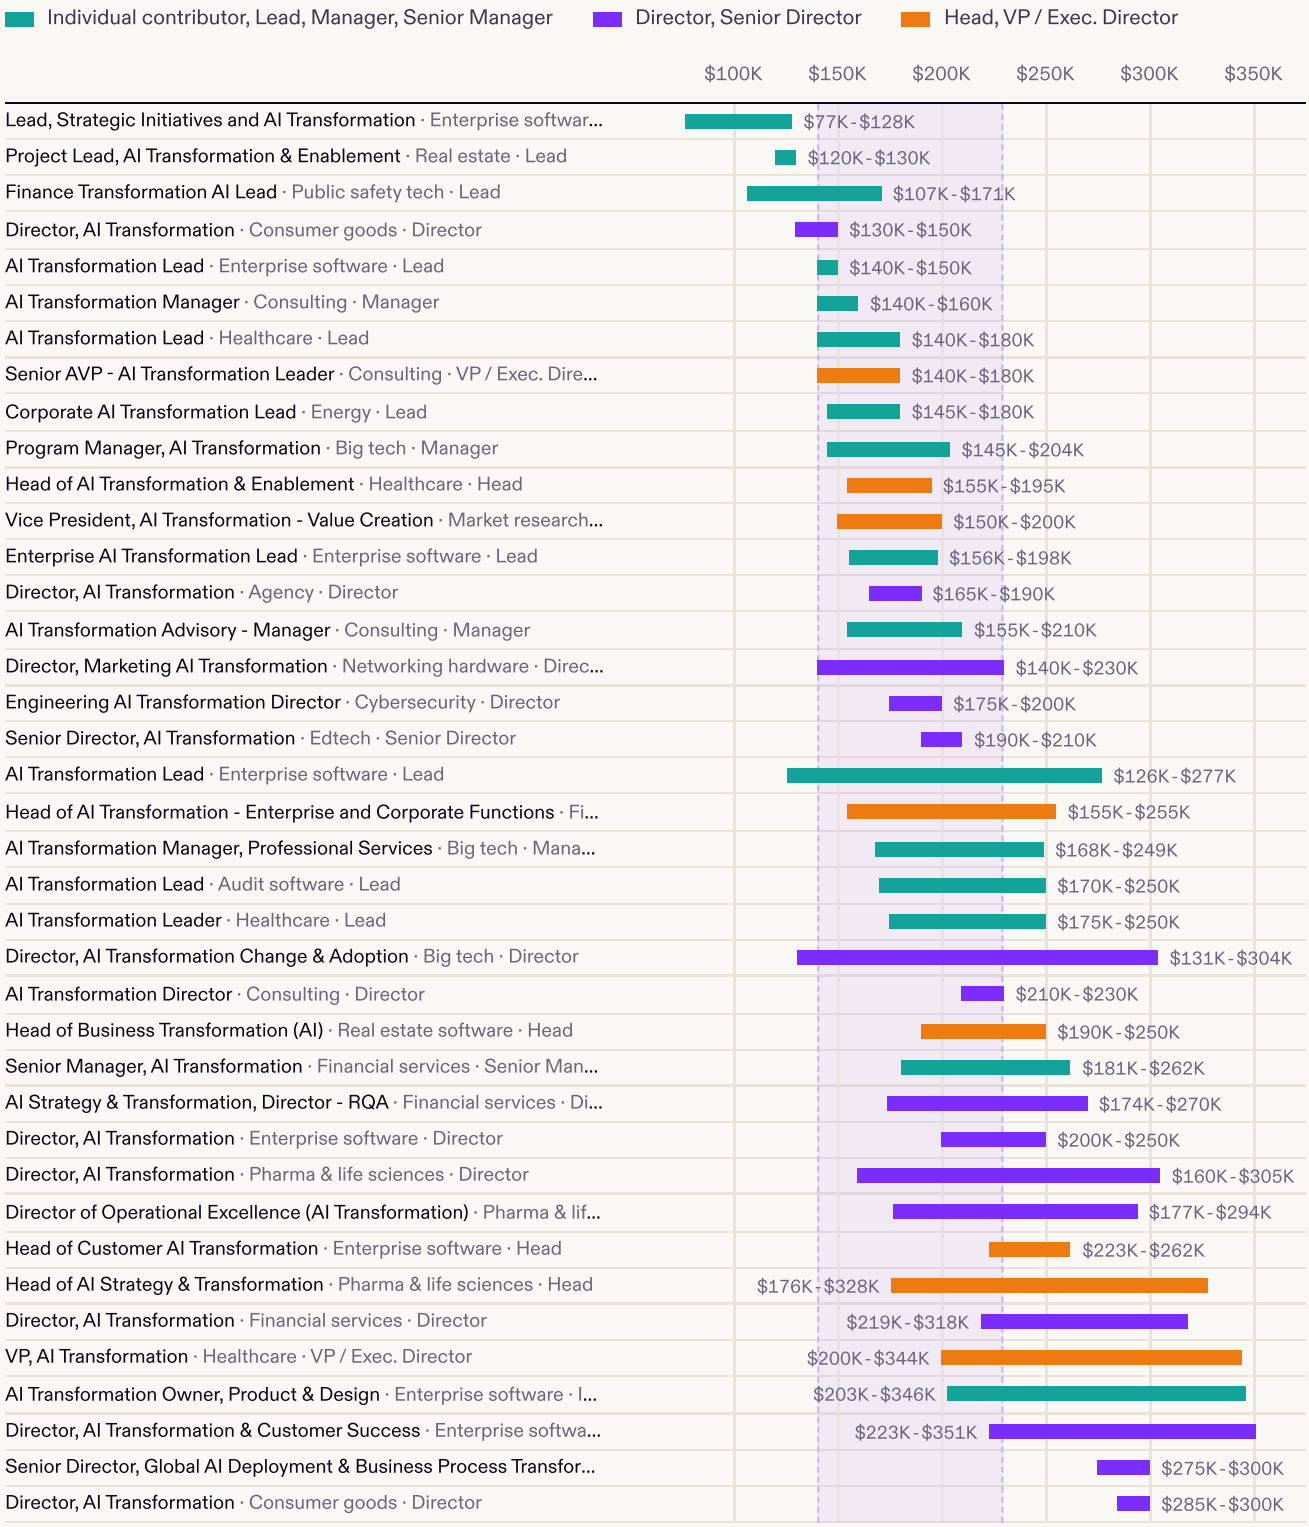
Exhibit sources and qualifications

EXHIBIT	SAMPLE AND SOURCE	QUALIFICATION
1 **The AI journey**	Melt framework; examples from practitioner interviews, 2026	A framework, not a measured maturity scale. Melt assigned stages from what interviewees described; not self-reported.
2 **Selecting AI initiatives**	Practitioner interviews, 2026	Approaches as described by the organizations using them. Listed, not ranked.
3 **Fintech onboarding**	One company; practitioner interviews, 2026	Results as described by the company.
4 **Themes in postings**	94 of 102 postings with a usable description	Public job postings, July-October 2026.
5 **Roles by seniority**	102 postings	Only roles whose title or description centers on AI transformation.
6 **Salary ranges**	39 of 102 postings that disclose pay	Base salary as advertised; excludes bonus and equity; locations differ.
7 **Operating model**	Practitioner interviews and job postings, 2026	A synthesis of several organizations, not a universal blueprint.
8 **Role orientation**	100 of 102 postings that could be classified	A different measure from the 72% in Exhibit 4.

EXHIBIT 6

All 39 disclosed salary ranges, by job title and industry.

Base salary range by job title and industry, $, postings that disclose pay (n = 39 of 102); shaded band indicates $140K-$230K



Ranges as advertised at the time of the scan, July-October 2026. Base salary only; ranges reflect different locations and exclude bonus and equity. Bar color shows the level stated in each posting.

Source: Melt analysis of public job postings, July-October 2026

ABOUT MELT

Melt is building the decision layer for AI in complex enterprise operations, from manufacturing and supply chain to service delivery.

Every day, decisions about customers, suppliers, resources and priorities determine whether a business runs to plan or loses time and money. Yet the context needed to make those decisions is scattered across systems, spreadsheets, emails and employees' experience. As a result, issues take too long to resolve, costly problems recur, and critical expertise remains trapped with individuals.

Melt reconstructs how operational decisions happen and captures the know-how behind them, so it stays with the organization rather than with individuals. It identifies where intervention can make the biggest difference and the measurable business impact an agent is expected to deliver, then deploys AI agents to help teams act before small issues become costly delays. Because ROI guides the process from the start, the focus stays on business value, not automation for its own sake. We are turning the operational know-how that runs complex businesses into an AI workforce that can scale across the enterprise.

About the founders. Melt's founders have built and led enterprise technology businesses with deep roots in manufacturing and industrial environments, from founding an industrial cybersecurity company and leading it through its acquisition by a global technology company to running enterprise cybersecurity, cloud and data products and working closely with manufacturing organizations. They combine that experience with a practical understanding of the gap between how complex businesses are designed to operate and how work actually gets done on the ground.

FROM RESEARCH TO ACTION

The knowledge gap	Melt turns the experience behind everyday operational decisions into knowledge the organization can use.
The value gap	Melt links AI initiatives to the business problems they are meant to solve and the measurable impact they should deliver.
The ownership gap	Melt connects agents to the people accountable for their performance and business outcomes.

