

AIGW ENTERPRISE WORKFORCE TRANSITION DEEP DIVE

When AI-Native Redesign Outran Agent Readiness

Meta's Project OT and the Organisational Limits
of AI-Led Workforce Restructuring

October 2026

EVIDENCE REVIEWED THROUGH

13 October 2026

CASE CONTEXT

United States · Technology and Digital
Services

CONTENTS

Contents

Executive Snapshot	03	7. The reversal was partial, not an exit from AI	11
At a glance	04	8. What the evidence says about human authority	12
1. Why this case matters	05	9. What Meta teaches about measuring an AI workforce transition	13
2. The plan, the pilots and the implemented organization	06	10. AIGW interpretation: the architecture moved unevenly	14
3. Role redesign: from specialist to builder	07	11. Implications for enterprise leadership and policymakers	15
4. The workforce reduction: what is established, and what is not	08	12. Evidence boundaries and conclusion	16
5. When more code did not translate directly into proportional product delivery	09	Notes and original sources	17
6. Employee trust became part of the operating problem	10		

META · DEEP DIVE

Executive Snapshot

In 2026, Meta attempted to rethink not only how its employees use artificial intelligence but how work, management and staffing should be organized when AI agents become part of production. Internally, the initiative was known as Project OT—Organization Transformation. Its design included smaller AI-supported product teams, more fluid specialist roles, ambitious restructuring scenarios, and movements of employees into new priorities. But its implementation soon encountered a combination of operational uncertainty, workforce resistance and concern about whether agents were delivering sufficient reliable value. Meta ultimately canceled planning for a second broad restructuring wave while continuing several AI-related initiatives. [1](#)

This was not simply a technology rollout gone wrong, nor does the available evidence justify saying that artificial intelligence replaced thousands of Meta employees. The most consequential distinction is between **plans**, **implemented workplace changes**, **measured outcomes**, and **later reversals**. The reduction of approximately 8,000 positions in May 2026 is independently documented in Meta's securities filings. By contrast, much of the detailed evidence on pod structures, product-development metrics, employee sentiment and incident workloads

comes from Reuters' review of internal records and interviews with people familiar with the program. These are valuable investigative findings, but they are not equivalent to Meta publishing a complete controlled evaluation of its AI-native organization. [123](#)

The case reveals a management problem relevant far beyond Meta. A company may be able to create smaller teams using AI coding tools before it has worked out who will supervise their output, how specialists will contribute, who retains formal authority, what happens when agent actions go wrong, or how the organization will rebuild trust among people asked to work differently. Faster production of intermediate artifacts—such as code changes—also does not establish equally rapid delivery of dependable products or better outcomes for workers. [1](#)

AIGW interpretation: Meta illustrates the risk of **architectural pre-commitment**: redesigning jobs and structures around anticipated technological capability before the institutional system required to sustain that capability is demonstrably ready. This is a hypothesis grounded in the sequence of events, not an independently established causal verdict. It does not mean smaller AI-enabled teams are necessarily inferior, or that Meta has stopped investing in artificial intelligence.

CASE PROFILE

At a glance

FIGURE 01

Evidence synthesis

Project OT: observed and unresolved

Question	Evidence-grounded answer
What changed?	Selected engineering and research groups experimented with or implemented small AI-assisted pods, fluid “builder” responsibilities and flatter coordination structures. ^[1]
What was planned but not fully executed?	Some planning scenarios contemplated shrinking selected teams by as much as 60%; a second company-wide restructuring wave planned for later 2026 was canceled. ^[1]
What happened to staffing?	Meta's May reduction affected approximately 8,000 employees; Reuters separately reported large-scale transfers into new AI-related initiatives. ^{[2] [9]}
What was measured?	Internal accounts documented growth in code-change activity, slower growth in features shipped, operational incident burdens and declining employee sentiment. These measures cannot be collapsed into one AI productivity or causation statistic. ^[1]
Who retained authority?	Meta said people—not AI systems—made personnel ratings and promotion decisions. Detailed authorization, supervision and rollback rules for individual agents are not publicly established in the cited record. ^[1]
What remains unknown?	Pod-specific net productivity; the exact contribution of AI to job cuts; long-term redeployment, pay and career outcomes; sustained incident rates; and the durability of individual pod structures. ^[1]

Source: AIGW synthesis of the cited original sources.

Note: Plans, implemented changes, reported indicators and unanswered questions are distinct.

META · DEEP DIVE

1. Why this case matters

AI workforce commentary often moves too quickly between two opposite claims. In one account, AI lets a small number of highly capable employees do the work of much larger teams. In another, the same technology is primarily a driver of layoffs. Neither description is an adequate map of organizational transition. Technology can affect task execution, workflow, supervisory structure, staffing budgets and employee experience through different mechanisms and over different time horizons.

Meta is unusually instructive because elements of those mechanisms became visible within one enterprise. Reuters' August 2026 investigation assembled internal planning documents, organizational posts and recordings, along with interviews with more than 20 people knowledgeable about the transformation. Meta confirmed the existence of Project OT and some of its planning contours, while disputing simplistic interpretations of the restructuring. The result is a

rare, if still incomplete, picture of an organization attempting to move from AI adoption to AI-native design at speed. [1](#)

AIGW uses the term **workforce architecture** for the institutional arrangements that determine how tasks and workflows are allocated, how jobs and capability are organized, who retains authority, how careers and staffing change, and how outcomes are measured. The Meta evidence touches several of these domains at once. It shows why high model ambition alone is not proof that the surrounding institution can be reorganized safely or effectively.

The comparison is not between an old organization that was definitively optimal and a new model that was definitively defective. The question is whether the new operating structure had demonstrated enough effectiveness, legitimacy and control to support the level of workforce restructuring being contemplated.

META · DEEP DIVE

2. The plan, the pilots and the implemented organization

The story began before the formal 2026 restructuring. According to internal material Reuters reviewed, a Meta product executive announced an AI-oriented pilot in July 2025 with five small technology pods. Each included roughly two or three engineers and a designer, equipped with AI tools. Rather than relying on longer conventional planning cycles, the experiment sought to build prototypes in roughly four-week sprints. The initial change therefore concerned how product ideas could be explored and tested by a small team—not proof that a whole engineering organization could immediately operate with dramatically fewer employees. [1](#)

An internal playbook circulated later in 2025 described a more extensive approach. The familiar separation between designers, engineers and other product specialists would become more fluid. Team participants would be described as builders, with specialist capacity pooled across projects. Hierarchical management would give way to smaller pods receiving direction from a lead. Reuters reproduced material contrasting traditional teams of roughly 10–20 with proposed AI-native teams of roughly 3–5. These figures describe organizational archetypes in planning materials, not a verified census of all Meta teams. [1](#)

In January 2026, Project OT formalized the larger ambition. Executive scenarios explored AI-supported, smaller teams and, for certain groups, potential reductions in team size as large as 60%. This was **not** a plan to dismiss 60% of Meta's entire workforce. Meta expressly stated that the most drastic figures referred to some teams, that scenarios combined redeployments, closing open positions and cuts, and that not every scenario was expected to proceed. [1](#)

Crucially, some organizational change did proceed. By June 2026, at least 11 engineering and research units had implemented versions of the smaller pod model, according to internal announcements and people cited by Reuters. That is meaningful evidence of an implemented operating-model experiment. It is not evidence that all engineering groups adopted the same structure, that all new pods had identical leadership systems, or that each delivered a proven productivity gain. [1](#)

The distinction matters: **design proposal → local pilot → implemented subset → measured result** is an evidence sequence, not a set of interchangeable labels.

META · DEEP DIVE

3. Role redesign: from specialist to builder

The proposed builder model asked employees to work across boundaries that had previously defined specialist jobs. In principle, AI tools could help engineers prototype interfaces, give designers more direct control over working code, and let compact teams advance products without the same division of labor. Shared specialists would still provide expertise where needed. This could make work more adaptive; it could also move responsibilities onto employees who were not trained or authorized to own them. [1](#)

What changed in the publicly accessible evidence is principally **role organization** and **team coordination**, not demonstrated universal skill conversion. A new label for workers does not itself establish that an employee acquired all the technical, judgment and domain capabilities implied by that label. Nor does a faster prototype establish that a product can be maintained, tested and safely released without specialist input.

The management structure created a related challenge. Reuters reported that some teams experimented with a “village” approach in which senior organizational leads held personnel-management authority while day-to-day pod leads directed work without conventional managerial powers. An internal post by one affected employee

raised practical questions about management tools and formal performance responsibilities. That testimony supports a problem of role clarity; it does not prove that every manager was stripped of the ability to supervise. [1](#)

There was also a sensitive distinction involving artificial intelligence itself. One internal description referred to AI systems supporting elements of the personnel-management process. Meta responded that actual performance ratings and promotions were made by **people, not AI**. The responsible interpretation is therefore narrow: documentation suggests experimentation around management processes, while final personnel decision-making remained human according to the company. The evidence does not establish automated promotion or dismissal decisions. [1](#)

AIGW interpretation: The central risk is not role flexibility as such. It is a possible mismatch between **responsibility, authority** and **support infrastructure**. A pod lead may be expected to coordinate performance without receiving the tools or formal mandate of a manager; a specialist may be expected to deliver across disciplines without a verified development pathway. These are workforce-architecture questions before they are technology questions.

META · DEEP DIVE

4. The workforce reduction: what is established, and what is not

Contemporaneous Reuters reporting in March and April 2026 had already described proposals for major workforce cuts, including an initial May wave and possible later reductions. Those articles document what was expected *before* management canceled the later wave; they are not, by themselves, proof that every reported cut was implemented. [45](#)

In May 2026, Meta carried out a material headcount reduction. Its July 29 Form 10-Q states that the action affected approximately **8,000 employees** across its Family of Apps and Reality Labs operations and generated **\$1.18 billion** of severance expenses. The company reported **75,472 employees as of June 30**, but explicitly warned that the figure still included approximately 8,000 people affected by the May action, most of whom would leave the reported count by the end of the third quarter. The June workforce figure is therefore not a clean post-reduction denominator. [23](#)

The cuts and Project OT were connected in timing, but the causal story is more complex than “agents replaced 8,000 employees.” At an April 30 internal town hall, Zuckerberg said growing AI infrastructure investment required the company to reduce spending elsewhere, including personnel costs. He also said the workforce cuts were not the same as the move toward AI-native team structures. Reuters, meanwhile, documented planning in which workforce reductions, reduced vacancies and possible redeployments were considered alongside organizational redesign. The record supports

describing these as **co-occurring dimensions of Meta's AI-era restructuring**, not a directly measured one-for-one substitution from human labor into agents. [16](#)

The evidence also points to actual workforce movement, not cuts alone. A June memo reported by Reuters referred to roughly **7,000 employees transferred** to newer initiatives related to AI workflows. The August investigation quoted Meta's position that thousands had moved toward company priorities, including work on model-training data. These are direct, attributable descriptions of redeployment. They do not tell us how many transfers were voluntary, permanent, upward in grade, accompanied by pay gains, or successful after several months. [19](#)

This is why workforce effects need to be measured as flows. Headcount reductions, vacancies left unfilled, transfers, new-team entries, attrition and the reconfiguration of jobs are distinct mechanisms. The Meta case documents several, but it does not publish a reconciled person-by-person workforce-flow ledger.

A further plan was not carried out. Reuters reported that management called off planning for a second broad restructuring wave, initially considered for November 2026, just before the May layoffs. Meta confirmed this and said the second wave was canceled before a definitive tally of affected employees was reached. This is a significant partial reversal. It is not evidence of a permanent guarantee against future function-specific adjustments or all future layoffs. [1](#)

META · DEEP DIVE

5. When more code did not translate directly into proportional product delivery

At the center of Project OT was an expectation that AI tools and agents would change the economics of development. Small teams might explore more ideas and ship products faster because code creation, analysis and prototyping would require less human time. But a faster intermediate step does not necessarily mean the entire production system can safely accelerate. [1](#)

The August Reuters investigation quoted internal posts about contrasting indicators. Changes to Meta's internal software platforms and infrastructure reportedly rose **220% year over year**, while changes resulting in new or upgraded features reaching users increased **36%**. These numbers are not directly comparable measures of one task and do not show that AI caused only a 36% productivity improvement. They use different units and different positions in the development process. Their analytical significance is the gap between measures of **activity** and measures of **downstream delivery**. [1](#)

This does not mean that all additional code was wasted. Infrastructure changes can improve quality, capacity, security or maintainability without immediately appearing as visible user features. Nor does the difference isolate AI as the only driver of the change. Yet it offers a strong

measurement warning: a company can generate substantially more work artifacts without demonstrating that end users receive a commensurate increase in useful, reliable outputs.

The operational evidence creates a second warning. Reuters cited internal accounts reporting major technical and security incidents up **40%** from a year earlier and time spent handling incidents up **70%**. The reports linked these concerns to increased AI-generated code and aggressive agent action, including insufficiently checked behaviors. However, Reuters did not publish complete metric definitions, team-level denominators, controlled counterfactuals or independent audit records. The correct conclusion is therefore that **operational reliability concerns were reported during the AI coding surge**—not that an experimentally established 40% increase in incidents was exclusively caused by agents. [1](#)

For enterprise leadership, the net-benefit calculation is wider than code generation. It must consider review and test effort, exception management, rollback, security exposure, incident investigation and remediation. Until those costs are measured alongside throughput, a claimed productivity gain may remain incomplete.

META · DEEP DIVE

6. Employee trust became part of the operating problem

The transformation also affected how staff understood their relationship with the company. Reuters reported that Meta planned and installed software on computers used by U.S.-based employees to capture mouse and keyboard interactions for potential AI training. The initiative raised privacy and consent concerns in a workplace where employees were simultaneously hearing about staff reductions and models that might replicate some of their tasks. [7](#)

In May, employees organized protests against the tracking programme. Reuters documented workplace flyers and worker organizing efforts. Later, the company paused the programme; according to subsequent town-hall reporting, the company said employee activity data captured through the initiative had not yet been used to train AI models and that any renewed programme might use opt-in consent. This sequence must not be shortened to the unsupported claim that Meta trained replacement agents on secretly harvested employee keystrokes. [17810](#)

The investigation also cited a company Pulse survey in which favorable employee sentiment reportedly declined from **74% to 55%**. As with the

internal operational statistics, the survey is meaningful attributed evidence, not an independent audit or a causal estimate. The public sources cannot apportion that change among monitoring practices, restructuring uncertainty, work intensity, management messaging or other factors. [1](#)

The broader institutional issue is still consequential. When employees are asked to supply data or redesign work that may change their own jobs, they need credible explanations of purpose, consent boundaries, use restrictions, reporting channels and what employment transition support exists. Without these, a technically feasible training-data programme can lose legitimacy before its effectiveness is established.

AIGW interpretation: Worker trust should not be treated merely as a morale statistic after the technical design is complete. It is part of implementation architecture. Cooperation, learning, knowledge transfer and safe escalation become harder when the relationship between AI adoption and employment security is ambiguous.

META · DEEP DIVE

7. The reversal was partial, not an exit from AI

Meta's response to the resulting pressure was not to discontinue AI development. Reuters reported that executives paused the input-capture programme, allowed some employees to move back toward previous teams, adjusted how they discussed workforce stability, and re-examined unusually broad manager spans. A June memo obtained by Reuters acknowledged that management had made mistakes and would seek to provide greater organizational stability. [19](#)

At a July 2 internal town hall, Zuckerberg reportedly acknowledged that progress in AI agent capabilities had been slower than he expected, and that some organizational bets had not yet produced anticipated results. His expectation of improvement within the following months was a forecast, not independently verified outcome evidence. [10](#)

Meanwhile, Meta continued to invest heavily in AI products and infrastructure. Its August public

statement, *The Future is for Everyone*, advocates broad access to personal AI, human invention and individual empowerment. The statement is evidence of Meta's declared strategic philosophy, not a measurement of how Project OT changed workers' jobs. A subsequent October 12 model release likewise shows continuing technical development, but does not retroactively validate the earlier pod and workforce decisions. [1113](#)

These developments make it inaccurate to describe Project OT as the final failure of Meta's AI strategy. The more precise conclusion is that **a particular pace and configuration of workforce restructuring met constraints strong enough to prompt changes in execution and cancellation of a planned further wave**. Some pilots and redeployments continued; the durability and comparative performance of individual pods remain unresolved.

META · DEEP DIVE

8. What the evidence says about human authority

AI-native design introduces different authority questions at different layers of the company. The first concerns software agents: which actions can be taken autonomously, what review or testing is required, who can stop an agent, and who owns the consequences? Reuters-reported incidents make these questions material, but the retrieved public record does not publish Meta's complete action-level approval, rollback and accountability matrix for Project OT workflows. [1](#)

The second concerns the organization itself: who has authority over workload allocation, priorities, performance evaluation, hiring decisions and careers as traditional management layers are flattened? Meta's statement that final ratings and promotions remained human is an important boundary. It does not, however, resolve all practical questions about the role and tools of pod leads, the burden on organizational heads overseeing large spans of employees, or who is equipped to mentor and develop specialists working in broader builder roles. [1](#)

The third is procedural: who can consent to or challenge an employee-data programme? The tracking controversy illustrates that data governance and workforce governance meet when the data is generated by people's work practices. [78](#)

External AI governance literature supports asking these questions explicitly. NIST's AI Risk Management Framework calls for organizational accountability structures, clearly differentiated human responsibilities, oversight design and ongoing testing and monitoring. It is a benchmark for how enterprises might evaluate such systems; citing it does **not** imply Meta violated NIST guidance or that NIST has evaluated Project OT. [12](#)

A workforce redesign is therefore incomplete when new software responsibilities and changed human titles appear faster than the decision rights, training and accountability systems needed to manage them.

META · DEEP DIVE

9. What Meta teaches about measuring an AI workforce transition

At least four different dashboards were visible in this case, but they should not be compressed into a single notion of organizational success.

FIGURE 02

Evidence synthesis

Measurement layers and evidence limits

Measurement layer	Visible indicator	What it does not demonstrate
Workforce flow	Approximately 8,000 affected in the May reduction; transfers to new priorities	Agents directly displaced every affected worker, or that every transfer was an improved career outcome ^{[2] [9]}
Team architecture	AI-assisted pods established in a subset of units	Pods produced durable net value across the enterprise ^[1]
Technical activity	More code changes on internal systems	Equivalent increase in shipped features or reliable productivity ^[1]
Operational quality	Internal incident and firefighting metrics	A clean causal estimate of agent-specific losses ^[1]
Worker experience	Reported decline in favorable sentiment	How much each policy or technology caused the decline ^[1]
Governance	Human final personnel decisions stated by Meta	Detailed action-level agent approval and rollback architecture ^{[1] [12]}

Source: AIGW synthesis of the cited original sources.

Note: Rows are analytical distinctions, not a common numerical scale or causal estimate.

This is a useful case of **measurement-layer confusion**. Management can observe increased development activity while reliability resources are becoming stretched, and can report workforce movement without knowing whether jobs have become more satisfying, sustainable or accessible. The correct research response is not to discard operational indicators, but to demand better links between each claimed benefit and the outcome at the same layer of the system.

For workforce strategy, a complete transition account would eventually include more than the number of people cut or transferred. It would distinguish actual role content, destination role, new capability demonstrated, grade and compensation, manager responsibility, retention, productivity, incident burden and worker agency. These data are largely unavailable for Project OT in the public materials reviewed.

META · DEEP DIVE

10. AIGW interpretation: the architecture moved unevenly

The episode exposes five tensions that an enterprise cannot solve merely by buying more capable models.

First, task acceleration can be mistaken for workflow completion. Generating code is only one part of developing a service that is maintainable and safe to release. More output can increase verification demand rather than eliminate it.

Second, role flexibility can precede capability formation. A builder may take on broader tasks, but someone must establish when that worker has acquired the judgment and specialist support needed to do them reliably.

Third, flatter structures can obscure managerial accountability. Large spans and fluid pod leadership are not automatically illegitimate; they require clear authority, accessible performance mechanisms and adequate developmental supervision.

Fourth, workforce accounting can become distorted when plans are mistaken for outcomes. A modeled 60% reduction for selected teams, an

executed 8,000-person workforce cut, reassignments into AI initiatives and a canceled second wave describe different workforce mechanisms.

Fifth, organizational trust can constrain technical implementation. A staff-data collection programme may be operationally feasible but institutionally costly if purpose, consent, and employment consequences are poorly explained.

These are **AIGW interpretations**, not findings of a randomized experiment. The evidence supports identifying the architecture gaps as plausible contributors to implementation trouble. It does not establish how much each mechanism caused the management reversal, nor whether a different structure would have yielded better results.

The broader lesson is not that enterprise AI should remain a collection of small pilots forever. It is that organizations need a way to increase the ambition of workforce redesign in step with demonstrable operational reliability, human capability, authority clarity and employee legitimacy.

META · DEEP DIVE

11. Implications for enterprise leadership and policymakers

For CEOs and boards

An AI-native restructuring should have a workforce consequence statement that distinguishes technology aspiration from demonstrated operating readiness. Before committing to large structural reductions, leaders should know which workflows have actually changed; which productivity benefits survive verification, review and exception costs; and how affected workers will move or exit. Budget pressure from AI capital spending should be reported separately from task-substitution evidence.

For CHROs

Role transitions need explicit destinations and evidence of outcome quality. Reassigning employees to AI training-data work may be necessary for a new product strategy, but it is not automatically reskilling or advancement. Track destination roles, manager support, learning, compensation, retention and worker voice. A flatter organization is not inherently a better learning organization.

For CIOs and chief AI officers

Measure code, prototypes and agent actions alongside shipped quality, recovery time,

production incidents and operational ownership. Require defined privileges, thresholds, intervention rights and rollback pathways for consequential agent actions. An agent's capacity to act is distinct from its authority to act.

For managers and learning leaders

Broader builder roles require competency expectations, specialist access, structured review and feedback. Removing middle-management layers without deciding who coaches and evaluates employees can move organizational work into unrecognized gaps.

For policymakers and research institutions

Aggregate layoff totals do not by themselves identify which tasks were displaced by AI. Public understanding would improve if employers could distinguish reductions from capital-budget tradeoffs, new AI-enabled roles, transfers, non-backfill, and true task substitution. Employee-data capture for AI training also raises questions about consultation, consent, privacy and the governance of work-generated data. The evidence here does not constitute a jurisdiction-specific legal finding about Meta.

META · DEEP DIVE

12. Evidence boundaries and conclusion

This Deep Dive rests on **two kinds of evidence that should not be made equivalent**. Formal corporate filings establish the May 2026 reduction and its accounting treatment. The distinctive internal operating details come chiefly from original Reuters reporting based on documents, recordings and interviews not independently held or audited by AIGW. Meta's statements in that reporting provide important corroboration and counter-positions. Different Reuters stories across several months are not wholly independent replications of every underlying claim. [123](#)

The report does **not** establish that:

- Meta replaced approximately 8,000 employees with working AI agents;
- all teams were reduced by 60%, or even that all considered team reductions occurred;
- every AI-native pod failed or that Meta abandoned its AI strategy;
- differences between code changes and shipped features are a clean measure of AI's causal productivity impact;
- AI caused each reported incident or all deterioration in employee sentiment;

- promotion and ratings decisions were made by AI rather than people;
- transferred staff enjoyed better pay, greater proficiency or durable employment;
- Meta's financial results can be allocated causally to Project OT.

What can be said is more informative. Meta experimented with and implemented smaller AI-supported work units, explored consequential restructuring scenarios, made real staffing reductions and transfers, faced documented employee concern, and publicly acknowledged limits in the timing of agent capabilities and the execution of organizational change. The decision to cancel a further restructuring wave illustrates the possibility of revising organizational plans even while AI investment continues. [12910](#)

The case therefore adds an essential dimension to enterprise AI research: **workforce transformation is not just the problem of accelerating tasks or reducing costs; it is the problem of building an institution capable of sustaining the new division of labor**. Strong technical tools cannot substitute for clear human authority, measured workflow outcomes, development pathways and trust. Nor can an attractive organizational vision substitute for evidence that the new operating model works.

SOURCES

Notes and original sources

1. Reuters investigative report. *Katie Paul, with additional reporting by Jeff Horwitz. "Mark Zuckerberg had a bold plan to replace Meta staff with AI. Here's how it imploded."* 26 August 2026. Internal documents, interviews and Meta's response.
<https://www.reuters.com/investigations/mark-zuckerberg-had-bold-plan-replace-meta-staff-with-ai-heres-how-it-imploded-2026-08-26/>
2. Meta Platforms, Inc. *Form 10-Q, quarter ended 30 June 2026, filed 29 July 2026. Official SEC company filing; workforce reduction, severance and reporting definitions.*
<https://www.sec.gov/Archives/edgar/data/1326801/000162828026050705/meta-20260630.htm>
3. Meta Platforms, Inc. *"Meta Reports Second Quarter 2026 Results."* 29 July 2026. Official company release; reported headcount and financial context.
<https://investor.atmeta.com/investor-news/press-release-details/2026/Meta-Reports-Second-Quarter-2026-Results/>
4. Reuters. *"Meta planning sweeping layoffs as AI costs mount."* 14 March 2026. Earlier reported planning; not proof of implementation.
<https://www.reuters.com/business/world-at-work/meta-planning-sweeping-layoffs-ai-costs-mount-2026-03-14/>
5. Reuters. *"Meta targets May 20 for first wave of layoffs; additional cuts later in 2026."* 17 April 2026. Contemporaneous plan, later partly canceled.
<https://www.reuters.com/world/meta-targets-may-20-first-wave-layoffs-additional-cuts-later-2026-2026-04-17/>
6. Reuters. *"Meta CEO Zuckerberg blames layoffs on capital spending, won't rule out more job cuts."* 30 April 2026. CEO town-hall statements; competing attribution account.
<https://www.reuters.com/business/world-at-work/meta-ceo-attributes-layoffs-plan-capex-wont-rule-out-further-job-cuts-2026-04-30/>
7. Reuters. *"Meta to start capturing employee mouse movements, keystrokes for AI training data."* 21 April 2026, updated 22 April. Employee-data programme.
<https://www.reuters.com/sustainability/boards-policy-regulation/meta-start-capturing-employee-mouse-movements-keystrokes-ai-training-data-2026-04-21/>

SOURCES

-
8. Reuters. “Meta employees launch protest against mouse-tracking tech at US offices.” 12 May 2026, updated 13 May. Workers’ response.
<https://www.reuters.com/sustainability/society-equity/meta-us-employees-organize-protest-against-mouse-tracking-tech-2026-05-12/>
 9. Reuters. “Zuckerberg says Meta made ‘mistakes’ in AI workforce shift.” 12 June 2026. Memo and workforce-reassignment reporting.
<https://www.reuters.com/business/metazuckerberg-admits-mistakes-made-ai-transformation-2026-06-12/>
 10. Reuters. “Meta’s Zuckerberg says AI agent tech progressing slower than expected.” 2 July 2026. Recorded town-hall reporting, including data capture update.
<https://www.reuters.com/business/zuckerberg-says-ai-agent-development-going-slower-than-expected-2026-07-02/>
 11. Meta / Mark Zuckerberg. “The Future is for Everyone.” 10 August 2026. Public vision and strategic position; not Project OT outcome evidence.
<https://www.meta.com/thefutureisforeveryone/>
 12. U.S. National Institute of Standards and Technology. *AI Risk Management Framework — Core, 2023* (current online AI RMF 1.0 presentation). External framework for human responsibilities and monitoring; not a Meta-specific compliance finding.
<https://airc.nist.gov/airmf-resources/airmf/5-sec-core/>
 13. Reuters. “Meta releases Muse Wonder AI models to challenge Anthropic and OpenAI.” 12 October 2026. Subsequent product context, not evidence of Project OT workforce outcomes.
<https://www.reuters.com/business/meta-releases-muse-wonder-ai-models-challenge-anthropic-openai-2026-10-12/>

Source and method note

AIGW reviewed formal Meta disclosures, Meta’s public statement, original Reuters investigations and contemporaneous reporting, and the NIST governance benchmark. Because internal organizational documents, employee surveys and technical-incident logs were accessed by Reuters rather than released directly to AIGW, this report consistently attributes their findings to Reuters’ reporting. Corporate disclosures verify workforce reductions and financial information but do not independently validate AI-agent productivity or all details of Project OT. The analysis differentiates implemented activities, proposals, canceled plans, reported indicators, company responses and AIGW interpretation. Its purpose is to identify mechanisms and evidence gaps, not to estimate effects across the entire technology industry or establish a conclusive causal evaluation of Meta’s organizational redesign.