

The Language of AI Deployment: A Longitudinal Documentary Study of Nine Large Organisations, 2022 to 2026

Richard Foster-Fletcher
MKAI Research Programme
research@mkai.org
Working paper, July 2026

Abstract

Between January 2022 and April 2026, corporate leaders publicly described their use of artificial intelligence thousands of times. This report analyses 155 verified public statements from nine major organisations: Microsoft, IBM, Salesforce, Shopify, Duolingo, Cisco, Meta, Amazon, and JPMorgan Chase. Drawing on earnings calls, shareholder letters, regulatory filings, press releases, and executive interviews, each statement was verified verbatim and tracked chronologically.

Statements were categorized into five core registers: descriptor, expansion, substitution, walk-back, and operational. The findings show that:

- Expansion language dominates public messaging across almost every organisation.
- Claims of AI replacing human labour concentrated in 2025 and early 2026.
- Subsequent qualifications or walk-backs followed on timetables ranging from ten weeks to never.

Corporate communication does not move through fixed, linear stages. Instead, these five categories function as fluid registers that often coexist across different speakers and channels. Crucially, while organisations frequently soften or qualify their public language following pushback, internal AI deployment continues without interruption.

Keywords: artificial intelligence, corporate communication, executive discourse, workforce, organisational change, documentary methods

Contents

1. Introduction.....	3
2. Method.....	3
3. The coding scheme	6
4. Microsoft.....	6
5. IBM.....	8
6. Salesforce.....	10
7. Shopify.....	13
8. Duolingo	15
9. Cisco	17
10. Meta	20
11. Amazon.....	22
12. JPMorgan Chase	24
13. Cross-case synthesis.....	26
Appendix A. Coding tables.....	26
Appendix B. Source register	37

1. Introduction

Organisations speak about artificial intelligence in a language of their own choosing, on channels they control, at moments they select. That language is evidence. It records what an organisation wanted its investors, employees, customers, and regulators to believe about the technology at a specific point in time. Because these statements are dated, the record can be tracked chronologically as an evolving sequence rather than a static position.

This report examines that public record for nine major organisations across the fifty-two months from 1 January 2022 to 30 April 2026. The timeline opens ten months before the public release of ChatGPT, providing a pre-generative baseline, and closes in the spring of 2026, by which point every company in the study had made public claims about operational AI deployment. The accounts that follow trace each organisation's language through short, verified quotations, followed by a cross-case synthesis identifying where these patterns recur.

The approach is strictly documentary, recording what organisations said without inferring unstated internal actions. Every quotation was fetched from a reachable source, confirmed word for word, and logged with its URL and access date, with unverified statements excluded regardless of their prominence in secondary reporting.

The central question is whether an organisation's public AI narrative can shift dramatically while underlying working practices move much slower. The nine records below test how public claims compare with operational realities across the corporate landscape.

2. Method

2.1 Design and inclusion

The study covers nine major organisations: Microsoft, IBM, Salesforce, Shopify, Duolingo, Cisco, Meta, Amazon, and JPMorgan Chase. An organisation qualified where a named senior executive made at least five verifiable public statements regarding their AI programme across at least three calendar years within the window. All nine met this threshold.

Key executive voices and statement counts include:

- Satya Nadella (13)
- Arvind Krishna (12)
- Marc Benioff (10)
- Harley Finkelstein (8) and Tobi Lütke (5)
- Luis von Ahn (10)
- Chuck Robbins (9) and Jeetu Patel (5)
- Mark Zuckerberg (12)
- Andy Jassy (11)
- Jamie Dimon (13)

The corpus totals 155 verified statements: 18 each for Microsoft, Salesforce, Shopify, Cisco, Meta, and JPMorgan Chase; 16 each for IBM and Amazon; and 15 for Duolingo.

Admissible primary sources included press releases, earnings-call transcripts, regulatory filings, executive posts on LinkedIn or X, conference keynotes, and direct-quote media interviews. Analyst commentary, speculative coverage without direct quotes, and unpublished internal communications were excluded. Each captured passage (15 to 80 words) describes AI capability, deployment, workforce implications, or operational integration. Statements are logged by date made rather than source publication date. All sources were verified on 6 July 2026.

2.2 Verification

Verification ran in two distinct stages. At the corpus-build stage, a candidate quotation entered the log only after the page carrying it had been fetched and read, and the wording copied directly from that live page. Candidates whose sources would not fetch were discarded before logging, and each organisation's corpus file records up to three such discards alongside the specific reasons for exclusion. At the second stage, an independent verification pass re-fetched every URL in the corpus and confirmed each quotation verbatim against the live page, checking attribution and date consistency at the same time.

The removal count from the second pass, which this method treats as reportable in its own right, was zero for all nine organisations, with no wording corrections of substance. That figure should be read against the design rather than as a claim of absolute perfection: because verification was built into capture, unverifiable material was excluded before it could be counted. The second pass therefore measured the stability of the sourced record on re-fetch rather than the error rate of a looser first draft. One process failure is disclosed: the first verification agent assigned to Salesforce reported completion while its output file had never been written, so the Salesforce corpus was re-verified from scratch in a second, fully logged pass whose output was confirmed on disk. In addition, the study lead personally re-fetched and confirmed nine of the highest-stakes quotations, those carrying named-executive workforce claims, across seven of the nine organisations.

Several original outlets sit behind access restrictions, principally *Bloomberg*, *The Wall Street Journal*, and the *Financial Times*. Where an executive's exact words from a restricted original were reproduced verbatim by a fetchable outlet, the study logs the reproduction and explicitly marks the origin. The outlets relied on in this way are *Axios*, *Fortune*, *CNBC*, *Entrepreneur*, *Banking Dive*, *PYMNTS*, and *CIO Dive*. The same protocol covers executive posts on LinkedIn and X, which rarely fetch directly: *BetaKit*, *Entrepreneur*, *Mind the Product*, and *Fortune* supply verbatim reproductions of the Shopify and Duolingo memos and their follow-ups, with the original channel noted in every entry. Statements available only as video, or only as paraphrase, were dropped regardless of how widely they were reported. The corpus files name the most significant of these exclusions, including 2022 broadcast remarks by Arvind Krishna and a widely quoted January 2024 Duolingo spokesperson line available verbatim only behind *Bloomberg's* access controls.

Earnings-call quotations are reproduced exactly as rendered by *The Motley Fool's* published transcripts, which occasionally carry automatic-transcription artefacts. Quoted passages were selected from stretches free of visible garbling, and the corpus notes the caveat. One further editorial disclosure: a single quotation in the JPMorgan Chase corpus joins two verbatim, non-adjacent bullet

points from the same shareholder-letter section with a marked ellipsis; both halves were confirmed on the page, and the join is flagged in the corpus and preserved in the case study.

As a final gate, after the case studies were drafted, every quoted span in all nine cases was checked programmatically against its organisation's verified corpus, with typographic quoting normalised. The check covered 334 substantive quoted spans and returned zero failures.

2.3 Coding

Each verified statement received one of five descriptive codes: Descriptor, Expansion, Substitution, Walk-back, and Operational, as defined in section 3. Coding was conducted in a single pass by a single coder working to explicit written decision rules, after verification was completed and without regard to any expected chronological shape.

The coding rules included specific tie-break principles:

- **Dominant frame:** The primary focus and register of the passage governs the code assignment.
- **Substitution vs Operational:** Substitution supersedes Operational whenever explicit labour reduction or direct workforce replacement is cited.
- **Walk-back criteria:** Walk-back requires a visible qualifying relation to an earlier, stronger claim or to public replacement anxiety.
- **Expansion vs Operational:** A financial, measured, and quantified register codes as Operational, whereas an aspirational or capability-driven register codes as Expansion. Revenue and order milestones for AI products code as Expansion, treating them as market demand claims, whereas explicit cost, productivity, and margin arithmetic codes as Operational.

Specific judgement calls are documented directly in the coding tables. These include Amazon's October 2025 workforce message, which was coded as Walk-back due to its reframing of an anticipated AI-driven reduction as general organizational leanness, with the Operational alternative noted. Similarly, Microsoft's July 2025 employee memo was coded as Walk-back for qualifying recent layoffs against overall stable headcount and rewarded internal talent. Conversely, Cisco's August 2024 restructuring statements were coded as Operational throughout, as the workforce impact was explicitly framed as a reallocation of expense to fund AI investment, with no statement claiming that AI was performing the affected work.

No inter-coder reliability statistic is reported because the coding was executed by a single researcher. To ensure complete methodological transparency, the full coding tables in Appendix A expose every individual assignment and decision basis to inspection.

2.4 Limitations

Four principal limitations bound the findings of this study.

First, the research reports corporate language, and public language is inherently curated. Corporate communication channels convey what organisations wish to communicate publicly, meaning the record itself cannot arbitrate between mere rhetoric and operational reality.

Second, the reproduction protocol places a small subset of quotations one step away from their original source. While every such step is explicitly marked in the source register, the chain of provenance is longer than a direct primary fetch.

Third, earnings-call transcript sources carry the unavoidable artefacts of automatic speech recognition. While this risk was mitigated by selecting passages free of visible garbling, transcription errors cannot be entirely eliminated.

Fourth, relying on a single coder trades inter-coder reliability statistics for absolute consistency of judgement. The explicit decision rules and statement-by-statement decision bases are published in full so that this methodological trade-off is completely visible.

Finally, the pilot case on Klarna, completed prior to this study, served as the structural template for the case study format. The nine cases examined here were drafted to the exact same specification regarding length, chronological structure, and quotation density.

3. The coding scheme

Five codes classify what each statement describes. They are descriptive, and the study does not assume they form a fixed sequence in time. Each coded statement sits on its organisation's timeline by date, and the order is whatever the dates produce.

Descriptor. AI mentioned as a general property of the organisation or product, with no specific capability claim.

Expansion. AI framed as enabling new capabilities, products, or experiences.

Substitution. AI framed as replacing or reducing human labour, or as delivering work previously done by employees.

Walk-back. Earlier substitution claims revised, qualified, or set beside statements about the continued importance of human work.

Operational. AI framed as operational efficiency, cost reduction, productivity, or margin, in more technical and less aspirational terms.

The brief for the study named no expected trajectory, and none was imposed. A record that stays flat, or that widens its claims across the window, is reported as plainly as one that narrows. Across the 155 coded statements the distribution is 9 Descriptor, 81 Expansion, 22 Substitution, 13 Walk-back, and 30 Operational.

4. Microsoft

Microsoft entered the window already bound to the fortunes of generative AI, through its partnership with OpenAI and through the Azure infrastructure on which OpenAI trained its models. Its public account of its own AI use between January 2022 and April 2026 comprises eighteen verified statements, drawn from four annual shareholder letters, six earnings calls, corporate announcements and one internal memo published on the company blog. Satya Nadella, chairman and chief executive, supplies thirteen of the eighteen, spread across all five calendar years, and sets the register throughout. Amy Hood, chief financial officer, supplies four, each in the idiom of capital expenditure, capacity

and margin. The remaining statement belongs to Scott Guthrie, then the executive vice president running the Cloud + AI group.

The 2022 baseline predates ChatGPT and reads accordingly. In March 2022, in the press release marking completion of the Nuance Communications acquisition, Guthrie described an acquisition that "brings together Nuance's best-in-class conversational AI and ambient intelligence with Microsoft's secure and trusted industry cloud offerings" and would "help providers offer more affordable, effective and accessible healthcare" (S1). AI appears here as a feature of an acquired product, with a named industry to serve. By late October the frame had widened: in the annual shareholder letter signed on 24 October, Nadella wrote that "we're seeing a paradigm shift as the world's large AI models become platforms themselves", and described customers applying those models to "writing assistance, code generation, and reasoning over data with our new Azure OpenAI Service" (S2). On the earnings call the next day he told an analyst that "the OpenAI partnership is a very critical partnership for us" and credited the company's own build directly: "That's what's driven, in fact, the progress OpenAI has been making" (S3). Five weeks before ChatGPT reached the public, Microsoft was already describing large models as a platform and its own cloud as the machinery beneath them.

The following year is the densest in the record. On 23 January 2023 the company announced a third, multibillion-dollar phase of its OpenAI partnership, which Nadella described as formed around "a shared ambition to responsibly advance cutting-edge AI research and democratize AI as a new technology platform" (S4). The product language arrived on 16 March, when the company blog introduced Microsoft 365 Copilot: "Today marks the next major step in the evolution of how we interact with computing, which will fundamentally change the way we work and unlock a new wave of productivity growth", Nadella said, calling natural language "the most universal interface" (S5). Hood's first appearance, in prepared remarks on the 25 July earnings call, ran at a cooler temperature. "Even with strong demand and a leadership position, growth from our AI services will be gradual as Azure AI scales and our Copilots reach general availability dates", she told investors, while committing to spend ahead of that revenue: "We expect capital expenditures to increase sequentially each quarter through the year as we scale to meet demand signals" (S6). The shareholder letter of 16 October then lifted the claim from one product to the whole industry. "This next generation of AI will reshape every software category and every business, including our own", Nadella wrote, setting the moment in the line of platform transitions the letter traced from PC and server through web and cloud (S7).

Through 2024 the two registers shared the same earnings calls. Opening the call of 25 April, Nadella reported a record quarter for the cloud business and described the Copilot line, from productivity software down to models and infrastructure, as "orchestrating a new era of AI transformation driving better business outcomes across every role and industry" (S8). Minutes later Hood gave the quarter its cost shape: "We expect capital expenditures to increase materially on a sequential basis", she said, and noted that "near-term AI demand is a bit higher than our available capacity" (S9). The shareholder letter of 18 October moved the object of the claims once more. Nadella now described "a rich tapestry of AI agents, which can take action on our behalf", able to "make small businesses more productive, make multinationals more competitive, make the public sector more efficient" (S10).

January 2025 brought the first public sizing of the AI business. On the 29 January earnings call Nadella reported that "Enterprises are beginning to move from proof of concepts to enterprisewide deployments" and that "our AI business has now surpassed an annual revenue run rate of \$13 billion, up 175% year over year" (S11). Hood folded the figure into her own measure, confirming that "our AI

business annual revenue run rate surpassed \$13 billion and was above expectations" in a quarter where "we also improved our operating leverage with higher-than-expected operating income growth" (S12).

The one break in the register came at midsummer. On 24 July 2025 Nadella published a memo to employees on the company blog, written in the wake of that year's job eliminations. "By every objective measure, Microsoft is thriving—our market performance, strategic positioning, and growth all point up and to the right", it began, before turning to what the metrics left out: "And yet, at the same time, we've undergone layoffs" (S13). Between those two sentences sat record capital spending and the observation that "Our overall headcount is relatively unchanged"; the memo gave the condition a name, "the enigma of success in an industry that has no franchise value" (S13). The quoted passage assigns the cuts no cause. Within the same document the mission language resumed: "What does empowerment look like in the era of AI?" Nadella asked, answering with a shift "from a software factory to an intelligence engine" (S14). By the shareholder letter of 15 October the sweep had returned in full. "More than any transformation before it, this generation of AI is radically changing every layer of the tech stack, and we are changing with it", he wrote, adding elsewhere in the letter that the shift "is reshaping not just our products and business models, but how we work" (S15).

By 2026 the claims had reached the scale of the economy. "We are in the beginning phases of AI diffusion and its broad GDP impact", Nadella told the earnings call of 28 January, adding that "we have built an AI business that is larger than some of our biggest franchises that took decades to build" (S16). Three months later, on the call of 29 April, the figure had almost tripled from the level first disclosed fifteen months earlier: "Our AI business surpassed \$37 billion ARR, up 123%", he reported, calling the moment the beginning of "one of the most consequential platform shifts that will change the entire tech stack as agents proliferate and become the dominant workload" (S17). On the same call Hood delivered the record's second and final workforce statement, a single sentence inside remarks that traced margins and expenses to continued investment in AI infrastructure and AI talent: "Total company headcount declined year over year as we focus on building high-performing teams that operate with pace and agility" (S18).

Read end to end, the record is a widening series of claims set against a steady financial undertone. The chief executive's object grows from an acquired healthcare product to a productivity suite, to every software business including Microsoft's own, to autonomous agents, and at last to gross domestic product, with the disclosed revenue figures growing in step. The chief financial officer prices the same phenomenon throughout, in capital expenditure, capacity, margin and, at the end, headcount. Across four years and four months, neither voice connects the technology to any reduction in jobs. The 2025 layoffs appear as an enigma set beside thriving metrics, and the 2026 headcount decline occupies one sentence about pace and agility, spoken on the call that sized the AI business at \$37 billion. Whatever joins the diffusion to the decline, the record leaves the joining unsaid.

5. IBM

IBM speaks about AI from both sides of the transaction. The company sells the systems it describes, and its accounts of internal use double as demonstrations for prospective buyers, a dual position that shapes nearly every entry in the record. Sixteen verified statements span April 2022 to April 2026, drawn from IBM Newsroom releases, quarterly earnings calls, a chairman's letter filed with the SEC,

and interviews given to business media. Twelve of the sixteen belong to Arvind Krishna, the chairman and chief executive, whose voice carries the record across all four years; the rest come from the senior executives Ric Lewis, Tom Rosamilia, Rob Thomas and Jim Kavanaugh.

The 2022 baseline is a hardware and survey story, told months before generative AI reset the vocabulary. On 5 April 2022, announcing the z16 mainframe and its on-chip AI accelerator through a Newsroom release, Ric Lewis, senior vice president for IBM Systems, said that "our clients can increase decision velocity with inferencing right where their mission critical data lives" (S1). The release put the machine's capacity at 300 billion inference requests a day, with real-time fraud detection the leading example. On 19 May the company published its Global AI Adoption Index, and Tom Rosamilia, senior vice president for IBM Software, said in the accompanying release that organisations were "looking to AI to help them address skills and labor shortages", trends he read as pointing to "the growing role that AI is playing both within organizations but also in society" (S2). In this first framing, AI stood in for the workers that companies could not find.

The register changed in the spring of 2023. On 1 May, in a Bloomberg News interview, Krishna said of IBM's roughly 26,000 back-office workers, "I could easily see 30% of that getting replaced by AI and automation over a five-year period" (S3). Bloomberg's original page is access-restricted, so the study logs the wording from Axios, which reproduced it verbatim the next day and reported an accompanying plan to pause hiring for such roles, some 7,800 positions. It is the earliest hard replacement claim in the study's record, and it arrived with a five-year horizon. Eight days later, on 9 May, IBM launched watsonx, its generative AI platform for enterprises, and Krishna's language in the Newsroom release faced outward: "Foundation models make deploying AI significantly more scalable, affordable, and efficient", he said, and the platform was built "so that clients can be more than just users, they can become AI advantaged" (S4).

The qualification followed within four months, attached to a much larger argument. On 22 August 2023, in a CNBC interview whose wording Fortune carried the same day, Krishna said he had been "slightly misquoted" in coverage of the May remarks, the planned reduction now described as roles left unfilled as people departed (S5). In the same appearance he moved the case from one company's back office to the working population at large: "you need to get productivity. Otherwise, quality of life is going to fall. And A.I. is the only answer we got" (S5). Fortune also carried his estimate that companies "can get the same work done with fewer people" (S5). By the third-quarter earnings call on 25 October, the internal account had settled into the idiom it would keep, efficiency with a sales figure alongside. IBM was applying AI to client support, HR and IT, Krishna told analysts, "to automate a significant portion of tasks and improve the productivity of our people", and the quarter's generative AI "book of business" came "in the low hundreds of millions of dollars" (S6).

The 2024 record moves between earnings figures and platform strategy. On 24 January, in prepared remarks on the fourth-quarter 2023 call, Krishna reported that "demand continues to increase and our book of business in the fourth quarter is roughly double the third-quarter amount" (S7). In the question-and-answer session of the same call he qualified his own headline number, telling the analyst Toni Sacconaghi, "this is not all revenue in the quarter. I would just begin with that statement to set it straight", and framing the figure as a reading of momentum and client sentiment (S8). The chairman's letter in the 2023 annual report, filed with the SEC on 11 March, carried the internal claims to shareholders: "IBMers are also embracing watsonx", with examples that included "generating code up to 60% faster" (S9). At THINK 2024 on 21 May, announcing that the Granite models would be released as open source, Krishna reached back into company history in the Newsroom release: "We

want to use the power of open source to do with AI what was successfully done with Linux and OpenShift" (S10). The passage continued, "Open means choice" (S10).

The second replacement claim and its defence arrived inside a single conversation. In a Wall Street Journal interview given on 5 May 2025, Krishna said AI had taken over the work of several hundred human-resources staff through the company's AskHR agent while IBM hired more programmers and salespeople; the Journal's page is access-restricted, so the study logs the wording from Entrepreneur, which reproduced the sentence, brackets included, on 8 May: "Our total employment has actually gone up, because what [AI] does is it gives you more investment to put into other areas" (S11). Where the 2023 correction trailed its claim by sixteen weeks, the 2025 defence travelled with the disclosure itself. On 6 May, opening THINK 2025, Krishna's Newsroom release declared a phase complete: "The era of AI experimentation is over", with advantage now coming from "purpose-built AI integration that drives measurable business outcomes" (S12). On 20 October, in a Newsroom release announcing an inference partnership with Groq, Rob Thomas, senior vice president for software and chief commercial officer, kept to the production framing: enterprises have options "when they're experimenting", but "when they want to go into production, they must ensure complex workflows can be deployed successfully" (S13).

By 2026 the efficiency language had hardened into run-rate arithmetic. On the fourth-quarter 2025 earnings call on 28 January 2026, Krishna told analysts that "AI is also a powerful productivity driver for IBM" and set out the trajectory: "In 2023, we set out on a goal to achieve \$2 billion of productivity savings exiting 2024. And today, we are well ahead of that, exiting 2025 with \$4.5 billion of annual run rate savings" (S14). The same remarks put the cumulative generative AI book of business above \$12.5 billion and counted more than 20,000 IBMers using Project Bob, an internal AI development tool, with average productivity gains of 45 per cent. On the same call Jim Kavanaugh, the chief financial officer, reported that "clients are moving beyond experimentation", that "Our consulting generative AI book of business surpassed \$2 billion in the quarter", and that IBM was "expanding our impact through client zero", the company's term for offering its own deployment experience to clients (S15). On the first-quarter 2026 call on 22 April, the internal claim reached its widest form: "IBM Bob, our AI-based software development system, is now generally available. Our entire developer workforce is using Bob with average productivity gains of 45%" (S16). Elsewhere in the same remarks, AI in consulting was described as "both a growth driver and a productivity engine" (S16).

The shape of the record stays steady across its four years. Replacement claims surface in interviews, and the qualifications sit close behind them, four months behind the 30 per cent projection of 2023 and inside the very interview that disclosed the AskHR replacements in 2025, while the Newsroom's platform announcements run to their own calendar in between. What moves is the unit of account. In 2023 the company counted workers who might be replaced, and within a year it was counting automated tasks and a book of business; by 2026 the units were an annual run rate of \$4.5 billion and a developer workforce reporting average gains of 45 per cent. By Krishna's account, total employment rose through the same years in which the reported savings grew. The record carries both figures forward and leaves them unreconciled.

6. Salesforce

Salesforce, the San Francisco company built on customer relationship management software, entered the study period with an established AI brand, Einstein, and a habit of quantifying its output. The

verified record examined here runs from January 2022 to April 2026 and comprises eighteen public statements drawn from earnings calls, corporate press releases, keynote and podcast remarks as reported at the time, a results announcement filed with the SEC, and the company's own newsroom. Marc Benioff, co-founder and chief executive, supplies ten of the eighteen statements across four years, and the record's changes of register are mostly his. The remaining voices, five senior executives and the corporate newsroom itself, either support his framings or tidy up after them.

The baseline year, 2022, speaks in the language of prediction. On the earnings call of 1 March 2022, covering fourth-quarter and full-year results, in prepared remarks transcribed by The Motley Fool, the source for each earnings-call quotation in this account, Gavin Patterson, president and chief revenue officer, described the customer Moderna as "intent on becoming the first fully digital biotech company, embedding analytics, AI, and automation across every step of their value chain", with Einstein and Tableau CRM there to let it "use predictive analytics to make better decisions" (S1). AI in this telling is a component inside a customer's operations, priced into the platform. By the autumn the claims had acquired a superlative. The newsroom release announcing Salesforce Genie at Dreamforce on 20 September 2022 counted the installed intelligence: "Einstein, which generates over 175 billion predictions every day, can now deliver personalization and predictions based on real-time data" (S2). David Schmaier, president and chief product officer, called Genie "our most significant innovation ever on the Salesforce Platform", one that "makes every part of Customer 360 more automated, intelligent, and real time" (S3). Ten weeks remained before ChatGPT reached the public, and the operative nouns were still predictions and automation.

The generative turn is dated 7 March 2023, when a press release presented Einstein GPT as the first generative AI for CRM and carried Benioff's framing sentence: "The world is experiencing one of the most profound technological shifts with the rise of real-time technologies and generative AI" (S4). The older accounting sat directly alongside, the same release stating that Einstein "currently delivers more than 200 billion AI-powered predictions per day" (S4). By the earnings call of 30 August 2023 the shift had become a market forecast. "AI, data, CRM, trust, let me tell you, we are at the dawn of an AI revolution," Benioff told analysts, describing "a new innovation cycle" that was "sparking a massive tech-buying cycle over the coming years" (S5). Brian Millham, president and chief operating officer, supplied the deployment evidence on the same call: "our Einstein-powered chatbot is handling 40,000 customer service sessions per month", with the customer PenFed resolving "20% of their cases on first contact with Einstein-powered chatbots" and doing so "with incredible ROI" (S6). Less than a fortnight later, on the Dreamforce keynote stage on 12 September 2023, Benioff delivered the record's sharpest capability caveat, in remarks reported at the time by ITPro: "They call them hallucinations. I call them lies. These LLMs are very convincing liars. They can turn very toxic, very quickly" (S7). The keynote cast generative AI as needing a "trust revolution", and the caveat was aimed at rival large language models (S7).

The 2024 record opens in September with a new product and a larger claim. A press release of 12 September 2024 unveiled Agentforce, the company's autonomous AI agents, ahead of that year's Dreamforce, and Benioff supplied the escalation: "Agentforce represents the Third Wave of AI—advancing beyond copilots to a new era of highly accurate, low-hallucination intelligent agents that actively drive customer success" (S8). The same release carried his ambition "to empower one billion agents with Agentforce by the end of 2025. This is what AI is meant to be" (S8). The systems he had called liars a year earlier were now advertised as accurate. On 3 December 2024, on the first earnings call after Agentforce went into production, the vocabulary moved from software to labour. "This is

the market for digital labor," Benioff said, and "Salesforce has become right out of the gate here, the largest supplier of digital labor" (S9). Millham placed the moment in company history: "we redefined enterprise software with cloud and social and mobile and predictive AI over the past 25 years and now we're doing it again with Agentforce" (S10). Amy Weaver, president and chief financial officer, described Salesforce as "customer zero" for its own agents, reported "a significant pickup in case deflections for our customer service organizations", and called it "just the beginning of a new world of humans and AI agents working together to fundamentally reshape the economics of business" (S11). A deflected case is one that no employee handles; the workforce subject entered the record through the finance function.

The explicit workforce claims arrived on 26 February 2025, in the question-and-answer portion of the earnings call, in response to a question about augmentation and replacement. "And we're not going to hire any new engineers this year," Benioff said, continuing: "We're seeing 30% productivity increase on engineering, and we're going to really continue to ride that up. And we're going to grow sales pretty dramatically this year" (S12). The same answer allowed that some 9,000 support agents might be "rebalanced" into sales and marketing (S12). He also repeated what he was telling his peers, that "we're the last generation of CEOs to only manage humans", since "every CEO going forward is going to manage, you know, humans and agents together" (S13). The call named the coming fiscal year "the year of digital labor" (S13).

What followed ran out of order. In an interview with *Fortune* published on 30 July 2025, Benioff questioned the layoff attributions then circulating among chief executives: "I keep looking around, talking to CEOs, asking, What AI are they using for these big layoffs? I think AI augments people, but I don't know if it necessarily replaces them" (S14). For this study, both that quotation and the remarks that follow were logged from a single *Fortune* article of 2 September 2025, fetched in its syndicated form on Yahoo Finance. The September remarks are the strongest replacement claim in the record. Speaking on *The Logan Bartlett Show* podcast, in comments *Fortune* reported that day with direct quotation, Benioff described his own support operation: "I've reduced it from 9,000 heads to about 5,000, because I need less heads" (S15). A year before, he explained, "there would be 9,000 people that you would be interacting with globally on our service cloud" (S15). On the same page the cut is framed as a rebalancing, "I was able to rebalance my headcount on my support", with the conversations now split evenly: "50% are with agents, 50% are with humans" (S15). The doubt about replacement was published on 30 July, and the description of a roughly four-thousand-head reduction followed five weeks later. In this record the softening precedes the claim it might have softened.

The final year is institutional in tone. The fourth-quarter and full-year fiscal 2026 results release of 25 February 2026, filed as an exhibit with the SEC, carried Benioff's account of a rebuilt company: "We've rebuilt Salesforce to become the operating system for the Agentic Enterprise, bringing humans and agents together on one trusted platform" (S16). The units of measure had multiplied: "Agentforce ARR reached \$800 million, up 169% year-over-year", and the statement went on to count "nearly 20 trillion tokens" converted into "more than 2.4 billion agentic work units to date, moments where AI wasn't just reasoning — it was delivering real work" (S16). On 29 April 2026 the company's newsroom published a feature titled **How Salesforce Is Reshaping Its Workforce in the Age of AI** and retold the support story in house style: "In early 2025, Salesforce deployed Agentforce to answer customer support questions on help.salesforce.com. A few months later, Agentforce was handling 2.6 million customer conversations with a 63% resolution rate and customer satisfaction scores that matched human agents" (S17). The feature reported that the company had "successfully redeployed

hundreds of support engineers" while "not backfilling other roles" (S17). Sanjeev Balakrishnan, executive vice president of customer success, supplied its moral: "When technology inevitably changes how we work — which is what happened when Salesforce adopted its own product so effectively — we need to step up and show our people that we will invest in their next chapter as heavily as we invest in the tech itself" (S18).

The record ends with the company retelling its own reduction as a redeployment. Across the four years the numbers changed dimension, from 175 billion daily predictions to four thousand fewer support heads, while the pitch of confidence stayed level. The one public expression of doubt about replacement appeared in July 2025, and the strongest replacement claim followed in September, reported on the same syndicated page. A reader can treat the July remark as a sincere view overtaken by events, or as public reassurance that the podcast arithmetic later exposed; the dates permit either reading and settle neither.

7. Shopify

Shopify, the Ottawa-based commerce company whose software runs the online stores of millions of merchants, has described its AI work through a small and stable cast. The verified record runs to eighteen public statements between January 2022 and April 2026, eight of them from Harley Finkelstein, the president who fronts the quarterly earnings calls, and five from Tobi Lütke, the founder and chief executive who writes the company's memos. The remaining five come from product executives, the chief financial officer and the corporate newsroom. The statements spread across earnings calls, newsroom posts, one podcast interview and two staff memos, and the memos, each tied to a decision about jobs, carry the heaviest sentences in the record.

The baseline year offers machine learning without an era attached. A layoff letter published in July 2022 contains no mention of AI at all, and the earnings call held the day afterwards treated machine learning as installed machinery. On that call, held on 27 July 2022 and transcribed by The Motley Fool, Finkelstein described Shopify Audiences as a tool that "helps merchants find new customers", in which "Machine learning algorithms build an audience of high-intent buyers tailored for that product, and the audience list is directly and securely exported to the merchant's ad network of choice" (S1). The fulfilment operation built around the Deliverr acquisition was presented in the same idiom: "Using software and machine learning", its hubs would "unpack, scan and inspect all inventory" and place stock "based on expected buyer demand" (S2).

The year's other voice was Lütke's, on The Knowledge Project podcast with Shane Parrish, in an episode published on 15 November 2022, two weeks before ChatGPT reached the public. Asked which technology excited him most, he answered: "Right now, you got to say machine learning; machine learning has been unreasonably effective already. But boy, do transformer models change everything" (S3). In the same conversation he described watching several fields "progressing at breakneck pace" and conceded that "my model of what's possible is falling behind" (S4). The transformer remark predates every AI product the company would later announce.

The era arrived in a layoff letter. On 4 May 2023 Lütke published a memo on the company newsroom announcing a 20 per cent workforce reduction and the sale of Shopify Logistics, the business the 2022 fulfilment remarks had described, to Flexport. The memo's only passage on AI reads: "But now we are at the dawn of the AI era and the new capabilities that are unlocked by that are unprecedented", and it goes on, with "now" italicised in the original: "A copilot for entrepreneurship is now possible.

Our main quest demands from us to build the best thing that is now possible, and that has just changed entirely" (S5). On the Q1 2023 earnings call the same day, Finkelstein described the Shop app's new assistant, launched that March, as "like having your own shopping concierge powered by ChatGPT, one of the most advanced technologies that we have seen in decades", and placed the company "in the early innings of unlocking the true power of AI" (S6).

Sidekick followed that summer as the copilot made product. Announcing Sidekick and Shopify Magic in its Editions Summer '23 showcase on 26 July 2023, the company's head of product for AI, Miqdad Jaffer, told TechCrunch by email: "There's still so much untapped potential for AI and entrepreneurs, and we're deeply committed to making the power of AI accessible to businesses of all sizes" (S7). The same interview claimed a responsibility "to keep the businesses we power on the cutting edge of technology", while the article recorded the guardrails then in force, with merchants reviewing generated copy before it went live and AI features unable to write to production systems (S7). On the Q2 2023 earnings call a week later, on 2 August, Finkelstein raised the register: "We believe AI is making the impossible possible, giving everyone superpowers to be more productive, more creative, and more successful than ever before" (S8). Sidekick, in the same prepared remarks, was "a commerce expert in their corner who is deeply competent, incredibly intelligent, and always available" (S8).

By early 2024 the claims carried an internal ledger. On the Q4 2023 call on 13 February 2024, Finkelstein said the tools shipped the previous year had "saved our merchants and our team thousands of hours of work, enabling us to ship faster and make great decisions quicker" (S9). The strategy, in the same remarks, involved "integrating AI at the heart of our platform" in a technological era "where AI becomes the most powerful sidekick for business creation" (S10). The register reached hiring in August 2024, when a newsroom release of 28 August announced Microsoft's former AI lead Mikhail Parakhin as chief technology officer, calling him "one of the finest machine learning (ML) crafters on the planet" and promising he would "push Shopify to the cutting edge, not just in ML and AI" (S11).

The March 2025 memo moved the language from products to conditions of employment. Dated 20 March 2025 and posted in full by Lütke to X on 7 April, it enters this record through the verbatim reproductions published by BetaKit and Entrepreneur. The memo set "reflexive AI usage" as a baseline expectation and stated that "Using AI effectively is now a fundamental expectation of everyone at Shopify", with AI use added to performance reviews (S12). On abstaining, Lütke wrote: "Frankly, I don't think it's feasible to opt out of learning the skill of applying AI in your craft; you are welcome to try, but I want to be honest I cannot see this working out today, and definitely not tomorrow" (S12). Before requesting headcount, teams "must demonstrate why they cannot get what they want done using AI" (S12). The memo predicted that "AI will totally change Shopify, our work, and the rest of our lives", credited staff with using AI "to get 100X the work done", and restated the company's task: "Our job is to figure out what entrepreneurship looks like in a world where AI is universally available" (S13).

The record's one retrospective arrived in December 2025. In a newsroom post of 10 December announcing the Winter '26 Edition, styled as the RenAIssance Edition, Vanessa Lee, vice-president of product, looked back at the launch: "In the summer of 2023, we shared our vision for Sidekick—a super-charged cofounder for every merchant" (S14). The account then conceded what no earlier statement had recorded: "Sidekick has gone from a tool with only minor usage to an essential part of the platform", the change credited to a new architecture shipped earlier in 2025 (S14). The post presented the Edition as "moving AI from hype to everyday help" (S14). Those three words, "only

minor usage", are the record's first public gauge of how the 2023 launch had been received, and they sit inside the sentence declaring the product essential (S14).

The 2026 statements pair the era language with financial results. In the results release of 11 February 2026, Finkelstein called 2025 "Shopify at full throttle" and described the year as one of "laying the rails for the new era of AI commerce" (S15). The same release gave the coming year its slogan: "2026 will be the year of the builders" (S15). Jeff Hoffmeister, the chief financial officer, made his single appearance in the same release: "With AI reshaping how buyers discover and purchase, we delivered these strong margins while investing in Catalog, Sidekick, Universal Commerce Protocol, and our full platform of commerce solutions" (S16). The release put the spending beside a 19 per cent quarterly free cash flow margin and 30 per cent annual revenue growth. On the earnings call the same day, Finkelstein told analysts that "AI is a resource multiplier" and that "the rules of what's possible are being rewritten in real time"; the call reported orders from AI search up fifteenfold since January 2025, "on a small base", and almost 4,000 custom apps generated through Sidekick in three weeks (S17).

On 24 March 2026 the newsroom announced that millions of merchants could sell inside ChatGPT, Microsoft Copilot, Google AI Mode and Gemini. Mani Fazeli, vice-president of product, set the posture in the announcement: "Agentic commerce isn't something we're reacting to; it's a vision we're bringing to life at the very frontier of commerce and AI" (S18). His coda made the ambition plain: "Wherever commerce goes next, our merchants will be there first" (S18).

Across the four years the direction of the language never reverses. The July 2022 letter said nothing about AI; ten months later the only AI passage in a second letter opened an era while a fifth of the staff departed. The sole measure of uptake anywhere in the record, three words in a relaunch post, arrived more than two years after the launch it described, folded inside the announcement of the product's second act. The strongest internal claims, the thousands of hours saved and the hundredfold output, remain without a public figure, and the language that made them has moved on to rails, multipliers and frontiers.

8. Duolingo

Duolingo, the Pittsburgh company behind the language-learning app of the same name, listed on Nasdaq in 2021 and had built its product around machine learning long before generative models became a public preoccupation. Its verified record for this study runs from March 2022 to April 2026 and comprises fifteen statements drawn from an SEC filing, the corporate blog, earnings calls, press releases, LinkedIn posts, a podcast appearance and two statements from unnamed spokespeople. Luis von Ahn, the co-founder and chief executive, speaks in ten of the fifteen across four years; a blog author, a learning-design director and the spokespeople supply the rest. The record is distinctive at both ends, with replacement framing that predates ChatGPT and a final year in which five consecutive statements qualify what came before.

The baseline is proprietary and confident in tone. In the annual report filed with the SEC on 4 March 2022, the company described using "data from over half a billion exercises completed daily to train sophisticated machine learning algorithms that we deploy to improve learning efficacy", including a personalisation model called BirdBrain "designed to predict the probability that a given learner will get an exercise right or wrong" (S1). Within a month the framing moved from personalisation to production. A post on the company blog dated 6 April 2022, written by Sophie Wodzack and asking

whether a standardised test could write itself, described GPT-3 generating items for the Duolingo English Test while test developers worked in what the post called a "human in the loop" process, "filtering, editing, and reviewing AI-generated content to produce test items that are indistinguishable from something written by actual humans" (S2). The post appeared seven months before ChatGPT's public release, earlier than any comparable description elsewhere in this study.

Von Ahn first speaks in the record on the fourth-quarter 2022 earnings call of 28 February 2023, in answers preserved in the Motley Fool transcript. On generative AI and content creation he assured analysts that "we're still going to have humans go over it", then described what had already changed: "it used to be the case that all of this was almost entirely generated by humans. By now, it just has to be reviewed by humans, and it's a lot quicker to review than to generate" (S3). Asked why the company stayed out of human tutoring, he set out the position plainly: "human tutors are great, but they're not getting any better. They're just there. Whereas technology, in particular AI, and in particular a lot of the stuff coming from generative AI, is just getting better and better every year. So, we're going to bet on that" (S4). Two weeks later the bet acquired its public ambition. A corporate press release of 14 March 2023 announced Duolingo Max, built on GPT-4, carrying a quote from von Ahn that opened "AI accelerates our mission to make high quality education available to everyone in the world" and went on to the claim that "I believe AI will allow us to eventually recreate the experience of a human tutor and scale it to everyone in the world" (S5).

The efficiency language acquired a workforce dimension at the start of 2024. On 9 January, TechCrunch reported that the company had cut roughly ten per cent of its contractors at the end of 2023; a spokesperson, disputing the label "layoffs", described GPT translating sentences while "human experts validate that the output quality is high enough for teaching and is in accordance with CEFR standards for what learners should be able to do at each CEFR level" (S6). Von Ahn took the subject up six weeks later, on the fourth-quarter 2023 earnings call of 28 February 2024, calling the press reports "kind of a trigger point for me" (S7). His answer disputed the characterisation while confirming the cause: "We did reduce our contractor force, but this was not like full-time employee layoffs or anything like that", he said, before naming "the use of AI" as probably the biggest reason for the reduction and stating a general rule, "if something can be done by AI, we're going to take the opportunity" (S7).

After a fourteen-month gap in the record, the rule became a named policy on 28 April 2025, when von Ahn issued an all-hands email declaring Duolingo AI-first and the company posted it to its LinkedIn page. Because the LinkedIn original does not fetch, the wording relied on here is the verbatim reproduction published by Mind the Product. The memo said "AI is already changing how work gets done" and continued "It's not a question of if or when. It's happening now" (S8). It announced a set of "constructive constraints", among them a gradual end to contractor use for work AI could handle, a role for AI in hiring and in performance reviews, and new headcount only where teams could not automate more of their work (S8). Two days later, on 30 April, a corporate press release announced 148 new language courses. "Developing our first 100 courses took about 12 years, and now, in about a year, we're able to create and launch nearly 150 new courses. This is a great example of how generative AI can directly benefit our learners", von Ahn said, elsewhere in the release crediting "the incredible impact of our AI and automation investments, which have allowed us to scale at unprecedented speed and quality" (S9). Jessie Becker, the senior director of learning design, noted in the same release that a single new course had once taken a small team years to build: "Now, by using

generative AI to create and validate content, we're able to focus our expertise where it's most impactful, ensuring every course meets Duolingo's rigorous quality standards" (S10).

The memo drew a public backlash, and the qualifications began within a month. In a follow-up LinkedIn post of 23 May 2025, logged here from the verbatim reproduction published by Entrepreneur because the original likewise does not fetch, von Ahn conceded "When I released my AI memo a few weeks ago, I didn't do that well" in providing clarity, and drew a line the memo had left undrawn: "To be clear: I do not see AI as replacing what our employees do (we are, in fact, continuing to hire at the same speed as before)" (S11). The same post kept the direction, saying AI "is going to fundamentally change the way we work, and we have to get ahead of it" (S11). By August the episode carried a number. Explaining on the second-quarter 2025 earnings call of 6 August why daily-active-user growth had come in towards the low end of guidance, an effect concentrated in the United States, von Ahn told analysts: "The reason we came towards the lower end was because I said some stuff about AI, and I didn't give enough context. Because of that, we got some backlash on social media. The most important thing is we wanted to make the sentiment on our social media positive" (S12).

The qualifications continued into 2026 with the adoption itself reaffirmed. Asked about becoming AI native on the fourth-quarter 2025 earnings call, held on 26 February 2026, von Ahn answered that "ultimately, we are a technology company. We have been really trying to adopt AI and technology as fast as possible. In fact, we've probably gotten in trouble for adopting AI fast, for example, by writing stupid memos that I wrote one time" (S13). In April the revision reached policy. Speaking on the Silicon Valley Girl podcast, in remarks logged here from Fortune's report of 13 April 2026 because the podcast channel does not fetch, von Ahn said the company had dropped AI usage from performance reviews. "At the end, we backtracked", he said, recalling the message to employees: "No. Look, the most important thing in your performance is that you are doing whatever your job is as well as possible. A lot of times AI can help you with that. But if it can't, I'm not going to force you to do that" (S14). In the same piece he added that "The reality is it's not yet the case that AI is better at coding than humans. I think you still really need engineers, and you're going to need them for a long time" (S14). A spokesperson's statement to Fortune the same day gave the reversal its corporate form, opening "Duolingo has used AI for years to personalize learning and expand access" before turning to the people: "Our teams' work depends on human judgment, expertise, and creativity. AI tools assist with that work; they don't make decisions or replace the people building Duolingo. What drives every decision we make is what's best for learners" (S15).

Read end to end, the record moves in one direction for three years and then spends its final year qualifying itself. The working arrangement described at either end barely differs, with humans filtering, editing and reviewing a model's output in April 2022 and AI tools assisting work that depends on human judgement in April 2026. Between those two descriptions sit a bet against human tutors, a contractor reduction attributed mainly to AI, a memo its author came to call stupid, and a user-growth cost quantified for investors. The company withdrew the performance-review requirement and softened the memo's framing while continuing the adoption that had prompted both, and no statement in the record closes the distance between the work as practised and the work as described.

9. Cisco

Cisco Systems, the San Jose company whose switches, routers and silicon carry a large share of the world's network traffic, described its use of AI in eighteen verified public statements between May 2022 and February 2026, drawn from newsroom releases, earnings calls and an acquisition announcement filed with the SEC. Chuck Robbins, chair and chief executive throughout the window, supplies ten of the eighteen, with at least one in every year. Jeetu Patel supplies five across four years, and his changing titles trace the subject's rise inside the company: executive vice president for security and collaboration in 2023, chief product officer by January 2025, president by the following June. Scott Herren, the chief financial officer, speaks twice, both times on the afternoon of 14 August 2024, the day Cisco announced a restructuring. One statement, an October 2022 product release, carries no named speaker at all.

The record opens at product level. In May 2022, in a corporate press release announcing a predictive-analytics engine for networking, Robbins said that "Our research for predictive networks has been tested and developed with customers, and early adopters are seeing major benefits saving them time and money", and that "The industry has been waiting for secure, proactive networking and only Cisco can do it right" (S1). On the Q4 FY2022 earnings call that August, AI sat inside a list of security investments, Robbins telling analysts that the company was "focusing on cloud-based offerings, best-in-class AI-driven threat detection, and end-to-end security architectures" (S2). In October 2022 an unsigned corporate release for WebexOne, framing such features as "audio intelligence", promised that "Webex's AI-powered background noise reduction technology will be available in Webex Contact Center and available on PSTN voice calls" (S3). Each claim from that year attaches to a working function inside a shipping product, and all three statements predate the public arrival of consumer chatbots at the end of 2022.

The register widened in 2023. In a March newsroom release announcing purpose-built AI features for Webex, Patel reached for the company's history, saying that "Cisco has decades of experience in AI with its industry-leading natural language understanding, and audio and video intelligence in Webex", and continued: "As we double down on our AI investment, we're empowering our customers to deliver exceptional hybrid work and customer experience outcomes based on their datasets" (S4). On the Q4 FY2023 earnings call on 16 August, Robbins moved AI to the centre of the growth case, telling investors that "the acceleration of AI will fundamentally change our world and create new growth drivers for us", adding that "AI has been an important element in our products for several years" (S5). The same call carried the first order milestone. Citing "early success in hyper-scale Ethernet AI fabric deployments", Robbins reported that "To date, we have taken orders for over half a billion dollars for AI Ethernet fabrics" (S6). Five weeks later, announcing the agreement to acquire Splunk for approximately \$28 billion in a release filed with the SEC on 21 September, he said the "combined capabilities will drive the next generation of AI-enabled security and observability" (S7).

April 2024 brought the widest claim in the record. In a newsroom release launching Hypershield, a security architecture billed as "AI-native", Patel said that "AI has the potential to empower the world's 8 billion people to have the same impact as 80 billion" (S8). Four months later the AI language met the workforce. On the Q4 FY2024 earnings call of 14 August, the day the restructuring plan was announced, Robbins framed the move around investment, telling analysts that "we remain laser-focused on growth and consistent execution as we invest to win in AI, cloud, and cybersecurity" and describing a plan intended "to allow us to both invest in key growth opportunities, as well as drive more efficiency in our business" (S9). Herren supplied the numbers in his prepared remarks: "we expect to impact approximately 7% of our global workforce with total estimated pre-tax charges of up

to \$1 billion" (S10). Asked in the Q&A whether the savings sat in the guidance, he moved the description away from cost, speaking of "finding efficiencies across the company so that we can pivot more resources, much like we did last year, into the fastest growth areas within the company" and asking analysts to "think of it more as reallocating versus being in pursuit of cost savings" (S11). Robbins added scale and motive, saying "we actually are shifting hundreds of millions of dollars into AI, into AI networking for cloud, into AI infrastructure, silicon, and cyber", because "the market's moving so quickly, we have to do that" (S12). In the company's account the reduction funds the investment, and no voice on the call describes AI as doing the work of the roughly seven per cent who would leave.

By 2025 the company was also selling protection for other organisations' AI adoption. Unveiling a product named AI Defense in a January newsroom release, Patel put the urgency in market terms: "In a dynamic landscape where competition is fierce, speed decides the winners" (S13). The Q3 FY2025 earnings release of 14 May kept the milestones moving, reporting AI infrastructure orders from webscale customers above \$600 million for the quarter and a \$1 billion full-year target passed a quarter early, while the quoted line from Robbins stayed general: "The momentum we are seeing with AI is fueled by the power of our secure networking portfolio, our trusted global partnerships, and the value we bring to our customers" (S14). In the release issued at Cisco Live in San Diego on 10 June, which described security being fused further into the network, the vocabulary turned agentic: "Safety and security are the defining challenges of the AI era—and agentic AI multiplies the risk, as every new agent is both a force multiplier and a fresh attack surface", Patel said (S15).

The final quarter of the window brings the external and the internal accounts into the same week. At Cisco Live EMEA in Amsterdam, in a release of 10 February 2026, Patel described software agents in occupational terms: "As agents take on critical enterprise roles, we're developing protections that work both ways: preventing agents from being compromised and controlling what they can access and do on our behalf" (S16). On the Q2 FY2026 earnings call the following day, Robbins reported that "AI infrastructure orders taken from hyperscalers totaled \$2.1 billion in Q2 compared to \$1.3 billion just last quarter and equal to the total orders taken in all of fiscal year '25", and on the same call he lifted the full-year expectation to more than \$5 billion in AI orders (S17). His prepared remarks then turned inward. "Today, the majority of our product developers are using AI coding assistant and working alongside agents, which help us innovate faster across our portfolio", he said, and "over 90% of customer experience support cases are touched by AI and automation, enabling us to resolve a greater proportion of complex cases within 1 day and contributing to our highest ever customer satisfaction scores" (S18). The following sentence of those remarks credited AI use across sales, security, the supply chain and corporate functions with cost savings and efficiency gains.

Read end to end, the record ascends without interruption or retraction; orders climb from half a billion dollars cumulative in August 2023 to \$2.1 billion in a single quarter of fiscal 2026, and the descriptive vocabulary travels from background noise reduction to agents in enterprise roles. Across four years and eighteen statements no speaker presents AI as a replacement for Cisco labour, and the one workforce reduction in the window is framed as a means of paying for AI. By the close Robbins reports most of the company's developers working alongside agents and nine in ten support cases touched by AI, with the savings now attributed to the technology, while headcount is still described in terms of reallocation. The corpus contains no statement that joins those two accounts, and the record ends with them side by side.

10. Meta

Meta Platforms entered the study window with its founder still in charge, and the record belongs mostly to his voice. Of the eighteen verified statements gathered between January 2022 and April 2026, twelve are Mark Zuckerberg's, spread across all five calendar years and delivered through earnings calls, letters to employees, open letters and one podcast appearance. Susan Li, the chief financial officer, supplies the financial counterpoint in two statements from the earnings record. The remaining four arrived through unbylined corporate channels, chiefly the AI at Meta blog and the company newsroom. Read in sequence, the statements move from infrastructure towards the composition of the workforce itself.

The baseline year belongs to infrastructure and recommendation. On 24 January 2022, ten months before ChatGPT, the AI at Meta blog announced the AI Research SuperCluster, a machine the company said "will be the fastest AI supercomputer in the world when it's fully built out in mid-2022" (S1). On the second-quarter earnings call of 27 July 2022, Zuckerberg gave the deployment figures: "about 15% of content in a person's Facebook feed and a little more than that of their Instagram feed is recommended by our AI, from people, groups, or accounts that you don't follow", adding that "we expect these numbers to more than double by the end of next year" (S2). The year ended in contraction. On 9 November 2022, in the letter to employees announcing roughly 11,000 job cuts, Zuckerberg named what the narrowing spared: "We've shifted more of our resources onto a smaller number of high priority growth areas — like our AI discovery engine, our ads and business platforms, and our long-term vision for the metaverse" (S3). The same letter described a company "focused on becoming even more efficient" (S3). It was the first document in the record to set an AI commitment inside a workforce reduction.

The language-model line opened the following February. A blog post of 24 February 2023 announced the public release of LLaMA, "a state-of-the-art foundational large language model designed to help researchers advance their work in this subfield of AI", under a gated research licence (S4). Eighteen days later, on 14 March 2023, Zuckerberg wrote to employees again, announcing around 10,000 further job cuts under the banner of a Year of Efficiency, with the AI passage once more inside the restructuring: "Our single largest investment is in advancing AI and building it into every one of our products. We have the infrastructure to do this at unprecedented scale and I think the experiences it enables will be amazing" (S5). The same letter mentioned "building AI tools to help engineers write better code faster" (S5), the record's first mention of AI applied to Meta's own engineering. In July the licence widened. The newsroom announcement of Llama 2 on 18 July 2023 made the models free for research and commercial use and stated the doctrine: "We believe an open approach is the right one for the development of today's AI models" (S6). "By making AI models available openly, they can benefit everyone", it continued (S6).

In 2024 the ambition acquired a new name. On the fourth-quarter 2023 earnings call of 1 February 2024, in remarks preserved in Meta's official investor-relations transcript, Zuckerberg told investors that "this next generation of services requires building full general intelligence" (S7). He presented the shift as a revision of his own expectations: "Previously I thought that because many of the tools were social, commerce, or maybe media-oriented that it might be possible to deliver these products by solving only a subset of AI's challenges" (S7). A newsroom announcement of 18 April 2024 supplied the consumer superlative, declaring that "we believe Meta AI is now the most intelligent AI assistant you can use for free", available across Meta's apps to "help you plan dinner based on what's in your

fridge" (S8). In an open letter published on the newsroom on 23 July 2024, accompanying a release the letter called "the first frontier-level open source AI model", Zuckerberg set out the trajectory: "This year, Llama 3 is competitive with the most advanced models and leading in some areas. Starting next year, we expect future Llama models to become the most advanced in the industry" (S9).

January 2025 brought the engineer claims. On The Joe Rogan Experience podcast that month, in a remark logged here from a fetchable Entrepreneur reproduction dated 20 January 2025, Zuckerberg said: "Probably in 2025, we at Meta, as well as the other companies that are basically working on this, are going to have an AI that can effectively be a sort of midlevel engineer that you have at your company that can write code" (S10). The same reproduction records him saying that a lot of the code in Meta's apps "is actually going to be built by AI engineers instead of people engineers" (S10). Nineteen days after the podcast version, in prepared remarks on the fourth-quarter 2024 earnings call of 29 January 2025, Zuckerberg restated it for investors: "I also expect that 2025 will be the year when it becomes possible to build an AI engineering agent that has coding and problem-solving abilities of around a good mid-level engineer. And this is going to be a profound milestone and potentially one of the most important innovations in history" (S12). In the question-and-answer session the agent was confined to internal development, "not going to be an external product anytime soon" (S12). The same remarks forecast the year "when a highly intelligent and personalized AI assistant reaches more than 1 billion people", with Meta AI expected "to be that leading AI assistant" (S11). Susan Li's outlook on the same call priced the programme: "We expect our full year 2025 capital expenses will be in the range of \$60 billion to \$65 billion", with the growth "driven by increased investment to support both our generative AI efforts and our core business" (S13).

Six months later the register lifted. On 30 July 2025 Zuckerberg published a letter at meta.com under the title Personal Superintelligence, and its central claim concerned self-improvement: "Over the last few months we have begun to see glimpses of our AI systems improving themselves. The improvement is slow for now, but undeniable. Developing superintelligence is now in sight" (S14). The letter distinguished Meta's personal vision from superintelligence "directed centrally towards automating all valuable work" and from AI as "a force focused on replacing large swaths of society" (S14). The executive who had promised investors a mid-level engineering agent in January was, by July, assigning the automation of valuable work to rival visions.

The workforce theme moved to the centre at the start of 2026. On the fourth-quarter 2025 earnings call of 28 January 2026, Zuckerberg told investors that "2026 is going to be the year that AI starts to dramatically change the way that we work" and declared that "We are now seeing a major AI acceleration" (S15). He described investment in "AI-native tooling so individuals at Meta can get more done" and a company "elevating individual contributors and flattening teams" (S15). Meta was, he said, "starting to see projects that used to require big teams now be accomplished by a single very talented person" (S15).

The distinguishing day of the case came three months later. On 29 April 2026, in prepared remarks on the first-quarter earnings call, Zuckerberg set himself against the replacement narrative: "My view of AI is very different from many others in the industry. I hear a lot of people out there talk about how AI is going to replace people. Instead, I think that AI is going to amplify people's ability to do what they want" (S16). He added that "People will be more important in the future, not less" (S16). The same prepared remarks carried the record's strongest team-shrinkage claim: "Before wrapping, I want to talk for a moment about how AI is transforming our work. We are seeing more and more examples where one or two people are building something in a week that would have previously taken dozens

of people months" (S17). The commitment that followed was to streamline teams "so they are not bigger than they need to be" (S17). Susan Li's segment on the same call supplied the arithmetic, attributing compensation growth to "technical hires we have added over the past year, particularly AI talent" before reporting: "We ended Q1 with over 77.9 thousand employees, down 1% from Q4 as the impact of headcount optimization efforts in certain functions was partially offset by hiring in priority areas of monetization and infrastructure" (S18).

Across the five years the capability claims rise without interruption, from the fastest supercomputer through the most intelligent free assistant to superintelligence in sight. The language about people moves less evenly. Each contraction of the workforce arrived with an enlargement of the AI commitment, and the two threads meet on 29 April 2026, when the rejection of replacement and the strongest account of shrinking teams shared one set of prepared remarks, the chief financial officer's arithmetic beside them. The company that says people will be more important is the company measuring how few of them a project now needs. The record keeps the two registers where they were spoken, side by side in a single transcript, and offers nothing that joins them.

11. Amazon

Amazon entered 2022 as a retailer, a logistics operator, and, through Amazon Web Services, the world's largest seller of cloud computing. Its public account of artificial intelligence runs through a small set of channels: annual letters to shareholders published on aboutamazon.com, quarterly earnings calls, product announcements from the corporate newsroom and the AWS blog, and, twice in the period, messages addressed to its employees. The principal voice is Andy Jassy, president and chief executive since mid-2021, who speaks in eleven of the sixteen statements gathered here and appears in every calendar year of the record, which runs from January 2022 to April 2026. Read in order, the statements move from machine learning as a component of particular products to artificial intelligence as the frame through which the company explains its capital spending, its margins, and the size of its workforce.

The baseline year treats the technology as an ingredient. Jassy's first letter to shareholders as chief executive, published on 14 April 2022, reached machine learning through Prime Video, recording what the company had learned about "producing compelling entertainment with memorable moments and using machine learning and other inventive technology to provide a superior-quality streaming experience" (S1). On 23 June, Jeff Barr, chief evangelist at AWS, announced the CodeWhisperer coding assistant in preview on the AWS News Blog: "Trained on billions of lines of code and powered by machine learning, CodeWhisperer has the same goal. Whether you are a student, a new developer, or an experienced professional, CodeWhisperer will help you to be more productive" (S2). On 10 November a corporate press release introduced the Sparrow warehouse robot, which, "Leveraging computer vision and artificial intelligence (AI)", could "recognize and handle millions of items", on a page that promised the machine "will take on repetitive tasks" and stated that "Robotics technology enables us to work smarter, not harder, to operate efficiently and safely" (S3). In each case the technology belongs to a named product doing a bounded job, with productivity, efficiency, and safety as the stated stakes.

The register changed within months of ChatGPT's public release. In the 2022 letter to shareholders, published on 13 April 2023, Jassy disclosed that "We have been working on our own LLMs for a while now", asserted that the technology "will transform and improve virtually every customer

experience", and committed the company to "invest substantially in these models across all of our consumer, seller, brand, and creator experiences" (S4). On the first-quarter earnings call two weeks later, on 27 April, he told analysts that "few folks appreciate how much new cloud business will happen over the next several years from the pending deluge of machine learning that's coming" (S5). When Amazon Bedrock reached general availability on 28 September, Swami Sivasubramanian, vice president of data and AI at AWS, said in the newsroom release that recent advances had "led to a surge of interest in generative AI, sparking new ideas that could transform entire industries and reimagine how work gets done" (S6). Transformation, in this year's usage, is something the technology will do to customer experiences and to other industries.

The 2024 statements attach the technology to the company's history and to its revenue. A corporate blog post of 1 February, bylined by two Amazon vice presidents, Rajiv Mehta and Trishul Chilimbi, launched Rufus in beta as "a generative AI-powered expert shopping assistant trained on Amazon's extensive product catalog, customer reviews, community Q&As, and information from across the web", on a page asserting that Amazon "has been using AI very expansively for 25+ years" and that generative AI would "change virtually all customer experiences" (S7). The 2023 letter to shareholders, published on 11 April 2024, set the technology on a historical axis: "Generative AI may be the largest technology transformation since the cloud (which itself, is still in the early stages), and perhaps since the Internet", describing a revolution that "will be built from the start on top of the cloud" (S8). By the third-quarter earnings call of 31 October the axis had figures beneath it, with Jassy reporting that "AWS's AI business is a multibillion-dollar revenue run rate business that continues to grow at a triple-digit year-over-year percentage", one "growing more than three times faster at this stage of its evolution as AWS itself grew, and we felt like AWS grew pretty quickly" (S9).

In 2025 the comparisons were put to work defending expenditure. Pressed on capital spending approaching one hundred billion dollars, Jassy told the fourth-quarter 2024 earnings call on 6 February that "AI represents, for sure, the biggest opportunity since cloud and probably the biggest technology shift and opportunity in business since the Internet", and that "virtually every application that we know of today is going to be reinvented with AI inside of it" (S10). The 2024 letter to shareholders, published on 10 April 2025, promised that "Generative AI is going to reinvent virtually every customer experience we know, and enable altogether new ones about which we've only fantasized", and called the moment "a once-in-a-lifetime reinvention of everything we know" (S11). Alongside the reinvention language the letter counted "more than 1,000 GenAI applications being built across Amazon" and gave the present a plainer description: "The early AI workloads being deployed focus on productivity and cost avoidance", and "This is saving companies a lot of money" (S11).

On 17 June 2025 the efficiency theme reached Amazon's own payroll. In a message to employees on generative AI, published in full on aboutamazon.com, Jassy wrote: "We will need fewer people doing some of the jobs that are being done today, and more people doing other types of jobs. It's hard to know exactly where this nets out over time, but in the next few years, we expect that this will reduce our total corporate workforce as we get efficiency gains from using AI extensively across the company" (S12). The same message counted "over 1,000 Generative AI services and applications in progress or built" and foresaw "billions" of AI agents (S12). Among the sixteen statements, this one alone draws the line in advance from AI to the size of the company's own workforce.

The reduction arrived on 28 October 2025, and the announcing voice changed. Beth Galetti, senior vice president of people experience and technology, told employees in a message on aboutamazon.com that approximately 14,000 corporate roles would be cut. In the passage carrying the announcement's reasoning, the technology comes first: "This generation of AI is the most transformative technology we've seen since the Internet, and it's enabling companies to innovate much faster than ever before (in existing market segments and altogether new ones)" (S13). The organisational claim follows directly: "We're convinced that we need to be organized more leanly, with fewer layers and more ownership, to move as quickly as possible for our customers and business" (S13). The message linked back to Jassy's June note while stopping short of its causal claim. Where Jassy had named efficiency gains from AI as the mechanism that would reduce the corporate workforce, Galetti gave leanness, layers, and ownership as the reasons for the reduction being made, with AI as the setting, the condition that made speed urgent.

The 2026 record speaks in agents and in arithmetic. On the fourth-quarter 2025 earnings call of 5 February 2026, Jassy told analysts that "the primary way companies will get value from AI is with agents", adding that "It's harder to build agents than it should be" (S14). The same remarks put AWS at a \$142 billion run rate and the chips business above \$10 billion, growing at "triple-digit percentages year over year" (S14). The 2025 letter to shareholders, published on 9 April 2026, measured the wave against the company's beginnings: "Three years after AWS launched commercially, it had a \$58 million revenue run rate. Three years into this AI wave, AWS's AI revenue run rate is over \$15 billion in Q1 2026 (nearly 260 times larger than AWS at that same point)—and ascending rapidly" (S15). The same letter defended capital expenditure of roughly two hundred billion dollars for the year. On the first-quarter call of 29 April, the last statement inside the study window, Jassy turned to the margin economics of custom silicon: "at scale, we expect Trainium will save us tens of billions of dollars of CapEx each year and provide several hundred basis points of operating margin advantage versus relying on other chips for inference" (S16).

Sixteen statements across fifty-two months trace an ascent that never reverses, beginning with machine learning as a feature of a streaming service and ending in run rates, agents, and basis points of margin on in-house chips. The two messages addressed to employees sit inside that ascent while answering to a different pressure. Jassy's June 2025 note said ahead of the event that efficiency gains from AI would reduce the corporate workforce. Galetti's October 2025 note, announcing the removal of roughly fourteen thousand roles, reached back to the June message while giving leanness and speed as its reasons. Four months separate the anticipation from the execution, and somewhere in those four months the causal claim dropped out of the company's account of itself.

12. JPMorgan Chase

JPMorgan Chase is the only organisation in this study whose 2022 baseline already reads like a finance department's account of AI. Jamie Dimon supplies thirteen of the eighteen verified statements, across all five calendar years, most of them from the annual shareholder letters that have carried his views on the firm's technology for two decades. Around him the file adds the global chief information officer, the president, the asset and wealth management chief executive, and the chief data and analytics officer. The record they produce moves in one direction. The workforce language hardens step by step across four years, and at no point does any of it soften.

The April 2022 letter, written before ChatGPT existed, frames AI as targeted spend with measurable returns. "Increasingly, we are investing more money (think hundreds of millions of dollars) each year on AI for very specific purposes," Dimon wrote, citing client insight generated from public data (S1). The same letter counted the yield: about \$100 million spent on fraud systems since 2017, fraud losses down 14% against volumes up almost 50%, and "about \$100 million in annual savings" attributed to the technology investments (S2). Other organisations in this study open the window with product launches or research claims. JPMorgan opens it with a savings figure.

The 2023 letter, published five months after ChatGPT's release, raised the register without abandoning the ledger. "Artificial intelligence (AI) is an extraordinary and groundbreaking technology," Dimon wrote, with "more than 300 AI use cases in production today" (S3). The same section imagined "new ways to augment and empower employees with AI," naming ChatGPT directly (S4). That October, in a Bloomberg TV interview whose wording this study logs from CNBC's verbatim reproduction because the original broadcast page is access-restricted, Dimon put the technology into generational terms: "Your children are going to live to 100 and not have cancer because of technology. ... And literally they'll probably be working 3½ days a week" (S5). The same report records that, asked whether AI would replace bank jobs, he said "of course" it would, that "technologies always replace jobs," and that displaced workers would move to new roles. The replacement acknowledgement and the redeployment promise arrive in the same breath, in October 2023, and that pairing never separates again anywhere in the file. A month later the register was regulatory, with chief information officer Lori Beer describing sessions "helping regulators understand how we build the generative AI models, how we control them, what are the new vectors of risk" (S6).

The 2024 letter is the record's rhetorical summit. "We are completely convinced the consequences will be extraordinary and possibly as transformational as some of the major technological inventions of the past several hundred years: Think the printing press, the steam engine, electricity, computing and the Internet, among others," Dimon wrote (S7). The use-case count had climbed past 400, with generative AI under exploration in "software engineering, customer service and operations, as well as in general employee productivity" (S8). And the workforce sentence took its first definite shape: "It may reduce certain job categories or roles, but it may create others as well. As we have in the past, we will aggressively retrain and redeploy our talent" (S9). At Investor Day that May, Mary Callahan Erdoes announced that "everyone coming in here will have prompt engineering training to get them ready for the AI of the future" (S10), and in September the president, Daniel Pinto, gave the cost version at a Barclays conference: "Either we'll be able to process a lot more for the same money or spend less" (S11).

Through 2025 the deployment story becomes internal infrastructure. Chief data and analytics officer Teresa Heitsenrether, quoted in a Harvard Business School case the company reviewed before publication, described the LLM Suite platform: "It looks like a ChatGPT window, but behind that, we link to our data and retain the option to swap in and out new models based on their strengths. Think of it as a fund of funds model" (S12). The firm's own technology blog recorded the scale in June: "the platform went from zero to 200,000 onboarded users within eight months" (S13). In October, back on Bloomberg TV in remarks carried verbatim by Fortune, Dimon supplied the cost-benefit line, "for \$2 billion of expense, we have about \$2 billion of benefits" (S14), and in the same interview moved the workforce language forward again: "I think people shouldn't put their head in the sand. ... It is going

to affect jobs. ... There will be jobs that it eliminates. ... For JPMorgan, if we're successful, we'll have more jobs—but there'll probably be less jobs in certain functions" (S15).

By 2026 the tense had changed. Speaking to investors in February, in remarks Fortune carried directly, Dimon confirmed the displacement as an accomplished fact: "We have displaced people from AI, and we offer them other jobs. They are usually well-trained and highly talented, very good at things" (S16). The April 2026 shareholder letter completed the sequence that the 2022 letter had begun. "The importance of AI is real — and while I hesitate to use the word transformational — it is," he wrote, judging that adoption "will likely be far faster than prior technological transformations, like electricity or the internet" (S17). And the conditional of 2024 became a certainty: "AI will definitely eliminate some jobs, while it enhances others. Our firm will have definitive plans on how we can support and redeploy our affected workforce" (S18). The two halves of that quotation come from separate bullet points in the same letter section, joined here as the corpus logs them, with an editorial ellipsis.

Set in date order, the same speaker's workforce sentence moves from "may reduce certain job categories" in April 2024, to "there will be jobs that it eliminates" in October 2025, to "we have displaced people from AI" in February 2026, to "will definitely eliminate some jobs" that April. Nothing in the file walks any of it back, and nothing needs to, because every version of the claim carries its redeployment clause inside it. The bank never made the promise that other organisations in this study later had to qualify; it conceded the displacement early, in small steps, each wrapped in retraining language, while the operational ledger ran continuously underneath, from \$100 million in fraud savings to \$2 billion of expense matched against \$2 billion of benefits. The file shows a firm that treated candour about job losses as a manageable disclosure rather than a crisis. Whether that steadiness reflects the discipline of a regulated institution or simply a chief executive who had priced the reaction in advance, the record does not say.

13. Cross-case synthesis

The nine records comprise 155 coded statements. Their distribution is 9 Descriptor, 81 Expansion, 22 Substitution, 13 Walk-back, and 30 Operational, and the table below gives each organisation's counts and its coded sequence in strict date order.

Organisation	D	E	S	W	O	Coded sequence by date
Microsoft	0	13	0	1	4	EEEEEOEEEOEEOWEEEEO
IBM	1	7	1	2	5	EDSEWOEOOEWEEOEO
Salesforce	2	7	5	2	2	DEEEEODESEOSSWSES W
Shopify	1	13	1	0	3	EEEDEEEEEOESEEEEEOE
Duolingo	1	2	5	6	1	DSSEWSSEOWWWW W
Cisco	2	12	0	0	4	EDEEEEEEOOOEEDEEO
Meta	0	10	4	1	3	EEOESESESOESWSO
Amazon	1	12	1	1	1	DEEEEEEEESWEE O
JPMorgan Chase	1	5	5	0	7	OOEESDEESOOOOOSSES

Primary findings

The first finding is the overwhelming weight of expansion. Eighty-one of the 155 statements frame AI as a new capability, product, or market opportunity, and Expansion is the most common code in seven of the nine records. The two exceptions are instructive. Duolingo's most common code is Walk-back, because its final year was spent qualifying its earlier language. JPMorgan Chase's dominant code is Operational, because its public record begins and ends in the balance sheet. Everywhere else, across four and a half years, the primary public register for corporate AI remained the announcement of what the technology would newly make possible.

The second finding addresses the clean sequence the five codes are often assumed to form: general mentions, capability claims, replacement claims, qualifications, and operational efficiency talk. As an ordered set of linear phases, that sequence appears nowhere in the nine records. No organisation moves through the five codes in order, and none divides its timeline into distinct stages of one code each.

What recurs instead is a smaller, reactive unit: the coupling of a substitution claim with a qualification, at time lags varying from minutes to never. IBM's May 2023 replacement figure met its correction ten weeks later, while its May 2025 disclosure that AI had taken over the work of several hundred HR staff arrived inside the same interview as its defence that total employment had actually risen. Amazon's June 2025 expectation of a smaller corporate workforce was met four months later by an execution message from a different executive that reached back to the note while declining its causal claim. Duolingo's qualifications began within a month of its AI-first memo and continued for eleven months across five consecutive coded statements, marking the longest sustained walk-back in the study.

Salesforce inverted the order entirely: Marc Benioff questioned whether AI replaces people in July 2025, seven weeks before stating he had cut support from 9,000 heads to about 5,000 because he needed fewer. Meta compressed the coupling to a single morning, explicitly rejecting the replacement narrative while simultaneously describing one or two people doing work that previously required dozens, both delivered in the same prepared remarks on 29 April 2026.

Three records break this coupling by never forming it. Microsoft's corpus contains no substitution statement at all; its only workforce statements set layoffs and headcount declines beside thriving metrics without explicitly joining the technology to the job cuts. Cisco's corpus contains neither substitution nor walk-back statements; its 7 per cent workforce action in August 2024 was framed as a reallocation of expense towards AI investment, with the CFO explicitly steering analysts away from reading it as cost reduction. JPMorgan Chase forms the replacement claim and never retreats from it. Its record demonstrates why no retreat was needed: from its initial acknowledgement in October 2023 that AI would replace jobs, every version of the claim carried its retraining and redeployment clause inside it. There was never an unguarded sentence to qualify. The walk-back, where it occurs across these cases, appears to be the price of an earlier statement that arrived without its institutional insurance.

Dynamics of role, channel, and timeline

Third, the codes sort by speaker as much as by chronological time. Chief financial officers account for eleven statements in the corpus, and ten of them code as Operational, with the eleventh coding as Expansion. No CFO statement codes as Substitution or Walk-back anywhere in the study. Amy Hood prices Microsoft's expansion in capex and operating leverage while Satya Nadella narrates platform

shifts; Susan Li translates Meta's superintelligence programme into expense growth and headcount optimisation; Scott Herren renames a workforce reduction as resource reallocation. Chief executives carry the broad aspiration and, where it exists, the replacement claim. The finance function speaks a parallel language in which AI appears as capacity, margin, and run-rate, regardless of the company's broader rhetorical posture. Reading any single organisation's record as a unified voice therefore misdescribes it. These registers coexist by executive role, and an organisation can occupy three of the five codes in a single earnings call.

Fourth, channel choice shapes the underlying claim. The sharpest numerical substitution claims in the corpus were spoken away from formal corporate channels: the 30 per cent back-office target to *Bloomberg*, the reduction from 9,000 to 5,000 support heads on a podcast, the mid-level engineer claim on a podcast, and job elimination comments on broadcast television. Annual shareholder letters and regulatory filings carry carefully hedged versions of the same claims, relying on words like *may*, *anticipate*, and *impact*. The clear exception is Amazon, which published the study's most explicit corporate substitution statement on its own newsroom site in June 2025 under the chief executive's byline. This exception demonstrates that channel dynamics represent a strong executive tendency rather than a fixed rule.

Fifth, the historical baseline was never empty. All nine organisations spoke publicly about AI or machine learning in 2022 prior to the release of ChatGPT, accounting for twenty-five statements (fifteen Expansion, six Descriptor, three Operational, and one Substitution). The single 2022 substitution statement belongs to Duolingo, whose corporate blog described GPT-3 writing test items with humans moved to filtering and review in April 2022, seven months before generative AI had a prominent public profile. The generative moment did not create the substitution frame; it amplified a narrative that organisations were already employing.

Substitution statements subsequently concentrated late in the observation window: one in 2022, four in 2023, three in 2024, nine in 2025, and five in the first four months of 2026. Walk-backs concentrated later still, with eleven of the thirteen instances falling in 2025 and 2026. Rather than neat, successive phases, the aggregate record reveals that replacement claims and public qualifications rose together toward the end of the timeline.

Finally, across every case examined, public qualification never slowed operational deployment. Every organisation that softened or walked back its language maintained, and in most cases accelerated, the internal adoption that the language described. Duolingo removed AI requirements from employee performance reviews while reaffirming its commitment to adopting technology as fast as possible. Salesforce's newsroom reframed support reductions as staff redeployment while simultaneously reporting that its agent platform revenue had grown 169 per cent. Meta's public assurances that people would be more important arrived alongside a rising capital expenditure budget and a shrinking overall headcount.

The walk-back, across this entire record, represents a revision of public description rather than operational practice. Internal deployment continues uninterrupted under adjusted wording. What the public record cannot resolve is whether organisations view their binding commitments in the external narrative or in the internal execution. The gap between those two positions remains the defining operational space.

Appendix A. Coding tables

One table per organisation. Codes: Descriptor, Expansion, Substitution, Walk-back, Operational. Coding was a single pass by one coder on 7 July 2026 under the decision rules in section 2.3; the Basis column records the decisive language for each assignment.

Microsoft

S#	Date	Speaker	Code	Basis
S1	2022-03-04	Scott Guthrie	Expansion	Nuance conversational AI enabling healthcare experiences
S2	2022-10-24	Satya Nadella	Expansion	large AI models become platforms; Azure OpenAI Service
S3	2022-10-25	Satya Nadella	Expansion	Azure supercomputing driving OpenAI progress
S4	2023-01-23	Satya Nadella	Expansion	multibillion OpenAI partnership; best AI infrastructure for builders
S5	2023-03-16	Satya Nadella	Expansion	Copilot will fundamentally change the way we work
S6	2023-07-25	Amy Hood	Operational	AI revenue growth gradual; capex pacing, financial register
S7	2023-10-16	Satya Nadella	Expansion	AI will reshape every software category including our own
S8	2024-04-25	Satya Nadella	Expansion	Copilot stack orchestrating new era of AI transformation
S9	2024-04-25	Amy Hood	Operational	capex up materially; demand above capacity, financial register
S10	2024-10-18	Satya Nadella	Expansion	AI agents acting on our behalf across work and life
S11	2025-01-29	Satya Nadella	Expansion	AI business \$13bn run rate; opportunity register
S12	2025-01-29	Amy Hood	Operational	operating leverage and margin paired with AI demand
S13	2025-07-24	Satya Nadella	Walk-back	layoffs set beside thriving metrics; enigma reframing
S14	2025-07-24	Satya Nadella	Expansion	software factory to intelligence engine mission reframe
S15	2025-10-15	Satya Nadella	Expansion	AI platform shift changing every layer of stack
S16	2026-01-28	Satya Nadella	Expansion	AI diffusion, GDP impact, TAM growth claim
S17	2026-04-29	Satya Nadella	Expansion	\$37bn ARR; agents become the dominant workload
S18	2026-04-29	Amy Hood	Operational	headcount declined framed as pace and agility

Counts: Descriptor 0, Expansion 13, Substitution 0, Walk-back 1, Operational 4.

Sequence: Expansion, Expansion, Expansion, Expansion, Expansion, Operational, Expansion, Expansion, Operational, Expansion, Expansion, Operational, Walk-back, Expansion, Expansion, Expansion, Expansion, Operational

Coding notes: S13 coded Walk-back on the qualifying relation to replacement anxiety (layoffs acknowledged and reframed against stable headcount and record AI investment); the quote itself attributes no cut to AI, and no prior Microsoft substitution claim exists in this corpus. S18 coded

Operational rather than Substitution because the quoted passage frames the headcount decline as team performance and expense discipline without an AI-replacement claim.

IBM

S#	Date	Speaker	Code	Basis
S1	2022-04-05	Ric Lewis	Expansion	z16 on-chip inferencing capability claim
S2	2022-05-19	Tom Rosamilia	Descriptor	survey framing; AI's growing role generally
S3	2023-05-01	Arvind Krishna	Substitution	30% of back-office replaced by AI and automation
S4	2023-05-09	Arvind Krishna	Expansion	watsonx launch; clients become AI advantaged
S5	2023-08-22	Arvind Krishna	Walk-back	slightly misquoted correction of May remarks; attrition framing
S6	2023-10-25	Arvind Krishna	Operational	internal task automation plus productivity of our people
S7	2024-01-24	Arvind Krishna	Expansion	book of business roughly doubled; demand register
S8	2024-01-24	Arvind Krishna	Operational	metric qualified as commitments not in-quarter revenue
S9	2024-03-11	Arvind Krishna	Operational	watsonx internally; code 60% faster; automate tasks
S10	2024-05-21	Arvind Krishna	Expansion	Granite open source; do for AI what Linux did
S11	2025-05-05	Arvind Krishna	Walk-back	total employment actually gone up after HR replacement
S12	2025-05-06	Arvind Krishna	Expansion	era of experimentation over; agent product push
S13	2025-10-20	Rob Thomas	Expansion	production-scale inference partnership claim
S14	2026-01-28	Arvind Krishna	Operational	\$4.5bn annual run-rate productivity savings
S15	2026-01-28	Jim Kavanaugh	Expansion	consulting GenAI book past \$2bn; largest quarter
S16	2026-04-22	Arvind Krishna	Operational	whole developer workforce on Bob; 45% productivity gains

Counts: Descriptor 1, Expansion 7, Substitution 1, Walk-back 2, Operational 5.

Sequence: Expansion, Descriptor, Substitution, Expansion, Walk-back, Operational, Expansion, Operational, Operational, Expansion, Walk-back, Expansion, Expansion, Operational, Expansion, Operational

Coding notes: S5 coded Walk-back because the logged note carries Krishna's explicit correction of the May 2023 coverage ("slightly misquoted", roles left open by attrition), which is the qualifying relation the code requires; the quoted passage alone would read as macro-productivity framing. S7 and S15 coded Expansion because book-of-business milestones are demand claims; S14 and S16 coded Operational because savings run-rates and internal productivity percentages are cost and margin arithmetic.

Salesforce

S#	Date	Speaker	Code	Basis
S1	2022-03-01	Gavin Patterson	Descriptor	Einstein predictive analytics as embedded customer-story property

S2	2022-09-20	Corporate press release	Expansion	Genie: can now deliver real-time predictions
S3	2022-09-20	David Schmaier	Expansion	most significant innovation ever on the platform
S4	2023-03-07	Marc Benioff	Expansion	Einstein GPT; profound technological shift claim
S5	2023-08-30	Marc Benioff	Expansion	dawn of an AI revolution; buying cycle
S6	2023-08-30	Brian Millham	Operational	chatbot session and resolution metrics with ROI
S7	2023-09-12	Marc Benioff	Descriptor	capability caveat; LLMs as convincing liars
S8	2024-09-12	Marc Benioff	Expansion	Agentforce as Third Wave of AI
S9	2024-12-03	Marc Benioff	Substitution	largest supplier of digital labor framing
S10	2024-12-03	Brian Millham	Expansion	redefining enterprise software again with Agentforce
S11	2024-12-03	Amy Weaver	Operational	case deflections; reshape the economics of business
S12	2025-02-26	Marc Benioff	Substitution	engineer hiring freeze on 30% AI productivity
S13	2025-02-26	Marc Benioff	Substitution	last generation of CEOs to only manage humans
S14	2025-07-30	Marc Benioff	Walk-back	AI augments people; questions layoff attribution
S15	2025-09-02	Marc Benioff	Substitution	9,000 heads to 5,000; I need less heads
S16	2026-02-25	Marc Benioff	Expansion	Agentic Enterprise platform; \$800m ARR milestone
S17	2026-04-29	Corporate newsroom	Substitution	2.6m conversations; parity with human agents
S18	2026-04-29	Sanjeev Balakrishnan	Walk-back	invest in people's next chapter; redeployment framing

Counts: Descriptor 2, Expansion 7, Substitution 5, Walk-back 2, Operational 2.

Sequence: Descriptor, Expansion, Expansion, Expansion, Expansion, Operational, Descriptor, Expansion, Substitution, Expansion, Operational, Substitution, Substitution, Walk-back, Substitution, Expansion, Substitution, Walk-back

Coding notes: S7 coded Descriptor as a general negative property claim about LLMs with no capability, workforce, or efficiency assertion of Salesforce's own. S9 and S13 coded Substitution because the digital-labour and humans-plus-agents framings present AI as workforce, which is the labour frame even when sold as opportunity. S11 coded Operational because no explicit labour reduction is claimed and the register is deflection metrics and business economics. Notable ordering fact carried to synthesis: the augmentation qualification (S14, July 2025) precedes the strongest substitution claim (S15, September 2025).

Shopify

S#	Date	Speaker	Code	Basis
S1	2022-07-27	Harley Finkelstein	Expansion	Audiences launch; ML builds high-intent buyer audiences
S2	2022-07-27	Harley Finkelstein	Expansion	ML-run fulfilment capability via Deliverr
S3	2022-11-15	Tobi Lütke	Expansion	transformer models change everything
S4	2022-11-15	Tobi Lütke	Descriptor	personal framing of ML pace; no company claim
S5	2023-05-04	Tobi Lütke	Expansion	dawn of AI era; copilot for entrepreneurship possible
S6	2023-05-04	Harley Finkelstein	Expansion	Shop assistant; early innings of AI power

S7	2023-07-26	Miqdad Jaffer	Expansion	untapped potential; AI accessible to all businesses
S8	2023-08-02	Harley Finkelstein	Expansion	AI superpowers; impossible made possible
S9	2024-02-13	Harley Finkelstein	Operational	thousands of hours of work saved internally
S10	2024-02-13	Harley Finkelstein	Expansion	AI at the heart of the platform strategy
S11	2024-08-28	Corporate press release	Expansion	CTO hire framed as ML capability acquisition
S12	2025-04-07	Tobi Lütke	Substitution	headcount conditioned on proof AI cannot do job
S13	2025-04-07	Tobi Lütke	Expansion	entrepreneurship in a world of universal AI
S14	2025-12-10	Vanessa Lee	Expansion	Sidekick now essential; concedes earlier minor usage
S15	2026-02-11	Harley Finkelstein	Expansion	laying rails for the era of AI commerce
S16	2026-02-11	Jeff Hoffmeister	Operational	strong margins while investing in AI
S17	2026-02-11	Harley Finkelstein	Operational	AI as leverage; less time on busy work
S18	2026-03-24	Mani Fazeli	Expansion	agentic storefronts at the commerce frontier

Counts: Descriptor 1, Expansion 13, Substitution 1, Walk-back 0, Operational 3.

Sequence: Expansion, Expansion, Expansion, Descriptor, Expansion, Expansion, Expansion, Expansion, Operational, Expansion, Expansion, Substitution, Expansion, Expansion, Expansion, Operational, Operational, Expansion

Coding notes: S5 coded Expansion because the quoted passage claims new capability; the surrounding memo cut roughly 20% of staff without attributing the cuts to AI, and that context is carried in the note rather than the code. S12 coded Substitution on the memo's headcount condition (teams must show AI cannot do the work before requesting people), which the logged note records. S14 contains a retrospective concession about the 2023 Sidekick launch ("only minor usage") but revises a capability claim rather than a substitution claim, so it stays Expansion under the scheme; the concession is material for the case study. No Walk-back appears in Shopify's coded record.

Duolingo

S#	Date	Speaker	Code	Basis
S1	2022-03-04	Corporate (10-K)	Descriptor	BirdBrain ML personalisation as standing product property
S2	2022-04-06	Corporate blog	Substitution	AI writes test items; humans shifted to review
S3	2023-02-28	Luis von Ahn	Substitution	content once human-generated now human-reviewed only
S4	2023-02-28	Luis von Ahn	Substitution	bet on AI over human tutors
S5	2023-03-14	Luis von Ahn	Expansion	Max launch; recreate a human tutor at scale
S6	2024-01-09	Company spokesperson	Walk-back	human validation reassurance amid contractor cuts
S7	2024-02-28	Luis von Ahn	Substitution	AI the biggest reason for contractor reduction
S8	2025-04-28	Luis von Ahn	Substitution	AI-first memo; work change happening now
S9	2025-04-30	Luis von Ahn	Expansion	148 courses in a year via generative AI
S10	2025-04-30	Jessie Becker	Operational	AI creates content; expertise focused where impactful
S11	2025-05-23	Luis von Ahn	Walk-back	I do not see AI as replacing employees
S12	2025-08-06	Luis von Ahn	Walk-back	backlash tied to DAU guidance; regret framing

S13	2026-02-26	Luis von Ahn	Walk-back	stupid memos self-deprecation while reaffirming adoption
S14	2026-04-13	Luis von Ahn	Walk-back	we backtracked on performance-review AI requirement
S15	2026-04-13	Company spokesperson	Walk-back	AI assists; does not replace the people

Counts: Descriptor 1, Expansion 2, Substitution 5, Walk-back 6, Operational 1.

Sequence: Descriptor, Substitution, Substitution, Substitution, Expansion, Walk-back, Substitution, Substitution, Expansion, Operational, Walk-back, Walk-back, Walk-back, Walk-back, Walk-back

Coding notes: S2 is a pre-ChatGPT substitution framing (GPT-3 item generation with humans moved to review), which places Duolingo's substitution language earlier than any other organisation in the study. S6 coded Walk-back as human-importance reassurance issued in direct response to contractor-cut coverage; S7, six weeks later, confirms the substitution plainly, so the walk-back and the confirmation coexist in the record. S12 and S13 revise the April 2025 memo episode rather than the underlying automation policy; S14 and S15 revise the policy itself. The record ends on five consecutive Walk-back codes.

Cisco

S#	Date	Speaker	Code	Basis
S1	2022-05-04	Chuck Robbins	Expansion	predictive networks; only Cisco can do it right
S2	2022-08-17	Chuck Robbins	Descriptor	AI-driven threat detection as portfolio investment focus
S3	2022-10-26	Corporate press release	Expansion	AI noise reduction arriving in Contact Center
S4	2023-03-28	Jeetu Patel	Expansion	decades of AI; doubling down on AI investment
S5	2023-08-16	Chuck Robbins	Expansion	AI will fundamentally change our world; growth drivers
S6	2023-08-16	Chuck Robbins	Expansion	\$500m AI Ethernet fabric orders milestone
S7	2023-09-21	Chuck Robbins	Expansion	Splunk; next generation of AI-enabled security
S8	2024-04-18	Jeetu Patel	Expansion	8 billion people with the impact of 80 billion
S9	2024-08-14	Chuck Robbins	Operational	restructuring framed as investment shift plus efficiency
S10	2024-08-14	Scott Herren	Operational	7% workforce impact as expense realignment
S11	2024-08-14	Scott Herren	Operational	cuts reframed as reallocation, not cost savings
S12	2024-08-14	Chuck Robbins	Expansion	hundreds of millions shifted into AI; market urgency
S13	2025-01-15	Jeetu Patel	Expansion	AI Defense to secure enterprise AI transformation
S14	2025-05-14	Chuck Robbins	Expansion	AI momentum; \$1bn order target beaten early
S15	2025-06-10	Jeetu Patel	Descriptor	agentic era risk framing; no capability claim
S16	2026-02-10	Jeetu Patel	Expansion	protections as agents take on enterprise roles
S17	2026-02-11	Chuck Robbins	Expansion	\$2.1bn quarterly AI infrastructure orders
S18	2026-02-11	Chuck Robbins	Operational	90% of support cases touched by AI; developer assistants

Counts: Descriptor 2, Expansion 12, Substitution 0, Walk-back 0, Operational 4.

Sequence: Expansion, Descriptor, Expansion, Expansion, Expansion, Expansion, Expansion, Expansion, Operational, Operational, Operational, Expansion, Expansion, Expansion, Descriptor, Expansion, Expansion, Operational

Coding notes: the August 2024 restructuring (S9, S10, S11) is coded Operational throughout because the 7% workforce impact is framed as expense realignment to fund AI investment; no statement claims AI performs the work of those affected, so Substitution does not apply under the scheme, and with no prior substitution claim in the record Walk-back does not apply either. Cisco's coded record contains no Substitution and no Walk-back across the full window.

Meta

S#	Date	Speaker	Code	Basis
S1	2022-01-24	Corporate blog	Expansion	fastest AI supercomputer claim (RSC)
S2	2022-07-27	Mark Zuckerberg	Expansion	AI feed share to more than double
S3	2022-11-09	Mark Zuckerberg	Operational	resources shifted to AI priorities within layoff letter
S4	2023-02-24	Corporate blog	Expansion	LLaMA release; state-of-the-art foundational model
S5	2023-03-14	Mark Zuckerberg	Expansion	AI the single largest investment; amazing experiences
S6	2023-07-18	Corporate press release	Expansion	open approach; models benefit everyone
S7	2024-02-01	Mark Zuckerberg	Expansion	full general intelligence now required
S8	2024-04-18	Corporate press release	Expansion	most intelligent free AI assistant claim
S9	2024-07-23	Mark Zuckerberg	Expansion	Llama to become most advanced in industry
S10	2025-01-20	Mark Zuckerberg	Substitution	AI as midlevel engineer writing code
S11	2025-01-29	Mark Zuckerberg	Expansion	leading assistant reaching one billion people
S12	2025-01-29	Mark Zuckerberg	Substitution	AI engineering agent at mid-level engineer ability
S13	2025-01-29	Susan Li	Operational	\$60-65bn capex driven by generative AI
S14	2025-07-30	Mark Zuckerberg	Expansion	self-improvement glimpsed; superintelligence in sight
S15	2026-01-28	Mark Zuckerberg	Substitution	single talented person replaces big teams
S16	2026-04-29	Mark Zuckerberg	Walk-back	AI amplifies people; against replacement narrative
S17	2026-04-29	Mark Zuckerberg	Substitution	one or two people do work of dozens
S18	2026-04-29	Susan Li	Operational	headcount optimisation beside premium AI hiring

Counts: Descriptor 0, Expansion 10, Substitution 4, Walk-back 1, Operational 3.

Sequence: Expansion, Expansion, Operational, Expansion, Expansion, Expansion, Expansion, Expansion, Expansion, Substitution, Expansion, Substitution, Operational, Expansion, Substitution, Walk-back, Substitution, Operational

Coding notes: S3 coded Operational on the same rule as Cisco's restructuring statements (resources realigned toward AI within a workforce action, with no claim that AI performs the work). S16 and S17 sit in the same prepared remarks on the same call: the explicit rejection of the replacement narrative and the one-or-two-people-do-the-work-of-dozens claim were delivered minutes apart, so Meta's walk-back coexists with its strongest substitution language rather than succeeding it.

Amazon

S#	Date	Speaker	Code	Basis
S1	2022-04-14	Andy Jassy	Descriptor	ML as ingredient of streaming quality
S2	2022-06-23	Jeff Barr	Expansion	CodeWhisperer preview; ML coding companion
S3	2022-11-10	Corporate press release	Expansion	Sparrow robot capability introduction
S4	2023-04-13	Andy Jassy	Expansion	own LLMs to transform every customer experience
S5	2023-04-27	Andy Jassy	Expansion	pending deluge of machine learning; new cloud business
S6	2023-09-28	Swami Sivasubramanian	Expansion	Bedrock GA; transform industries, reimagine work
S7	2024-02-01	Corporate blog	Expansion	Rufus generative shopping assistant launch
S8	2024-04-11	Andy Jassy	Expansion	largest transformation since cloud, perhaps Internet
S9	2024-10-31	Andy Jassy	Expansion	multibillion AI run rate; triple-digit growth
S10	2025-02-06	Andy Jassy	Expansion	biggest opportunity since cloud, defending capex
S11	2025-04-10	Andy Jassy	Expansion	reinvent every experience; cost-avoidance workloads noted
S12	2025-06-17	Andy Jassy	Substitution	fewer people in some jobs; corporate workforce reduction
S13	2025-10-28	Beth Galetti	Walk-back	cuts framed as leanness, not AI replacement
S14	2026-02-05	Andy Jassy	Expansion	agents the primary route to AI value
S15	2026-04-09	Andy Jassy	Expansion	\$15bn AI run rate versus early AWS
S16	2026-04-29	Andy Jassy	Operational	Trainium saves capex; operating margin advantage

Counts: Descriptor 1, Expansion 12, Substitution 1, Walk-back 1, Operational 1.

Sequence: Descriptor, Expansion, Expansion, Expansion, Expansion, Expansion, Expansion, Expansion, Expansion, Expansion, Substitution, Walk-back, Expansion, Expansion, Operational

Coding notes: S13 is a judgement call. Galetti announces the roughly 14,000-role reduction that Jassy's June message anticipated, yet frames it as leanness, layers, and ownership with AI as context rather than cause, and the message links back to the June note. That reads as a qualification of the earlier causal attribution at the moment of execution, so it codes Walk-back; the alternative reading (Operational, leanness as efficiency register) is noted for transparency. S3 codes Expansion on the robot capability claim; the same page's remark that Sparrow "will take on repetitive tasks" while employees focus elsewhere is carried in the note and belongs to the case account.

JPMorgan Chase

S#	Date	Speaker	Code	Basis
S1	2022-04-04	Jamie Dimon	Operational	targeted AI spend with identifiable returns
S2	2022-04-04	Jamie Dimon	Operational	fraud AI quantified at \$100m annual savings
S3	2023-04-04	Jamie Dimon	Expansion	extraordinary technology; 300 use cases; future success

S4	2023-04-04	Jamie Dimon	Expansion	augment employees with LLMs including ChatGPT
S5	2023-10-02	Jamie Dimon	Substitution	3.5-day week; jobs replaced, workers moved (note)
S6	2023-11-10	Lori Beer	Descriptor	regulator walk-through of model control; no capability claim
S7	2024-04-08	Jamie Dimon	Expansion	consequences as transformational as printing press, electricity
S8	2024-04-08	Jamie Dimon	Expansion	400 use cases; GenAI exploration across domains
S9	2024-04-08	Jamie Dimon	Substitution	may reduce certain job categories; retrain and redeploy
S10	2024-05-20	Mary Callahan Erdoes	Operational	prompt training for all joiners; hours saved
S11	2024-09-10	Daniel Pinto	Operational	process more for same money or spend less
S12	2025-04-28	Teresa Heitsenrether	Operational	LLM Suite architecture; fund-of-funds model swapping
S13	2025-06-03	Corporate tech blog	Operational	zero to 200,000 users in eight months
S14	2025-10-07	Jamie Dimon	Operational	\$2bn expense against \$2bn of benefits
S15	2025-10-07	Jamie Dimon	Substitution	jobs eliminated in certain functions; retrain ahead
S16	2026-02-24	Jamie Dimon	Substitution	we have displaced people from AI already
S17	2026-04-06	Jamie Dimon	Expansion	transformational; faster than electricity or internet
S18	2026-04-06	Jamie Dimon	Substitution	will definitely eliminate some jobs; redeployment plans

Counts: Descriptor 1, Expansion 5, Substitution 5, Walk-back 0, Operational 7.

Sequence: Operational, Operational, Expansion, Expansion, Substitution, Descriptor, Expansion, Expansion, Substitution, Operational, Operational, Operational, Operational, Operational, Substitution, Substitution, Expansion, Substitution

Coding notes: the 2022 baseline is already Operational (quantified savings), which no other organisation in the study shows. Substitution language hardens across the window within the same speaker: "may reduce certain job categories" (2024), "going to affect jobs... jobs that it eliminates" (2025), "we have displaced people from AI" (February 2026), "will definitely eliminate some jobs" (April 2026). No statement codes Walk-back: redeployment and retraining language accompanies each substitution claim from the first, rather than arriving afterwards as qualification.

Appendix B. Source register

Every statement in the corpus with its source URL. All URLs were fetched and the quoted wording confirmed on 6 July 2026, with an independent re-verification pass the same day. The companion workbook carries the full statement log, including the verbatim passages, verification status, and contextual notes.

Microsoft

S#	Date	Speaker	Channel	URL
S1	2022-03-04	Scott Guthrie, Executive Vice President, Cloud + AI Group	Corporate press release, completion of Nuance Communications acquisition (news.microsoft.com)	https://news.microsoft.com/source/2022/03/04/microsoft-completes-acquisition-of-nuance-ushering-in-new-era-of-outcomes-based-ai/
S2	2022-10-24	Satya Nadella, Chairman and CEO	Annual shareholder letter, Microsoft 2022 Annual Report (letter signed October 24, 2022)	https://www.microsoft.com/investor/reports/ar22/index.html
S3	2022-10-25	Satya Nadella, Chairman and CEO	Q1 FY2023 earnings call, Q&A response to Mark Murphy (J.P. Morgan); Motley Fool transcript	https://www.fool.com/earnings/call-transcripts/2022/10/25/microsoft-msft-q1-2023-earnings-call-transcript/
S4	2023-01-23	Satya Nadella, Chairman and CEO	Corporate press release/blog post announcing third phase of OpenAI partnership ("multiyear, multibillion dollar investment")	https://blogs.microsoft.com/blog/2023/01/23/microsoftandopenaiextendpartnership/
S5	2023-03-16	Satya Nadella, Chairman and CEO	Official Microsoft Blog post introducing Microsoft 365 Copilot (post authored by Jared Spataro; Nadella quoted)	https://blogs.microsoft.com/blog/2023/03/16/introducing-microsoft-365-copilot-your-copilot-for-work/
S6	2023-07-25	Amy Hood, Chief Financial Officer	Q4 FY2023 earnings call, prepared remarks (FY24 outlook); Motley Fool transcript	https://www.fool.com/earnings/call-transcripts/2023/07/25/microsoft-msft-q4-2023-earnings-call-transcript/
S7	2023-10-16	Satya Nadella, Chairman and CEO	Annual shareholder letter, Microsoft 2023 Annual Report (letter signed October 16, 2023)	https://www.microsoft.com/investor/reports/ar23/index.html
S8	2024-04-25	Satya Nadella, Chairman and CEO	Q3 FY2024 earnings call, opening prepared remarks; Motley Fool transcript	https://www.fool.com/earnings/call-transcripts/2024/04/25/microsoft-msft-q3-2024-earnings-call-transcript/
S9	2024-04-25	Amy Hood, Chief Financial Officer	Q3 FY2024 earnings call, prepared remarks (outlook); Motley Fool transcript	https://www.fool.com/earnings/call-transcripts/2024/04/25/microsoft-msft-q3-2024-earnings-call-transcript/
S10	2024-10-18	Satya Nadella, Chairman and CEO	Annual shareholder letter, Microsoft 2024 Annual Report (letter signed October 18, 2024)	https://www.microsoft.com/investor/reports/ar24/index.html
S11	2025-01-29	Satya Nadella, Chairman and CEO	Q2 FY2025 earnings call, opening prepared remarks; Motley Fool transcript	https://www.fool.com/earnings/call-transcripts/2025/01/29/microsoft-msft-q2-2025-earnings-call-transcript/

S12	2025-01-29	Amy Hood, Chief Financial Officer	Q2 FY2025 earnings call, prepared remarks; Motley Fool transcript	https://www.fool.com/earnings/call-transcripts/2025/01/29/microsoft-msft-q2-2025-earnings-call-transcript/
S13	2025-07-24	Satya Nadella, Chairman and CEO	CEO memo to employees, published on the Official Microsoft Blog ("Recommitting to our why, what, and how")	https://blogs.microsoft.com/blog/2025/07/24/recommitting-to-our-why-what-and-how/
S14	2025-07-24	Satya Nadella, Chairman and CEO	CEO memo to employees, published on the Official Microsoft Blog ("Recommitting to our why, what, and how")	https://blogs.microsoft.com/blog/2025/07/24/recommitting-to-our-why-what-and-how/
S15	2025-10-15	Satya Nadella, Chairman and CEO	Annual shareholder letter, Microsoft 2025 Annual Report (letter signed October 15, 2025)	https://www.microsoft.com/investor/reports/ar25/index.html
S16	2026-01-28	Satya Nadella, Chairman and CEO	Q2 FY2026 earnings call, opening prepared remarks; Motley Fool transcript	https://www.fool.com/earnings/call-transcripts/2026/01/28/microsoft-msft-q2-2026-earnings-call-transcript/
S17	2026-04-29	Satya Nadella, Chairman and CEO	Q3 FY2026 earnings call, opening prepared remarks; Motley Fool transcript	https://www.fool.com/earnings/call-transcripts/2026/04/29/microsoft-msft-q3-2026-earnings-transcript/
S18	2026-04-29	Amy Hood, Chief Financial Officer	Q3 FY2026 earnings call, prepared remarks; Motley Fool transcript	https://www.fool.com/earnings/call-transcripts/2026/04/29/microsoft-msft-q3-2026-earnings-transcript/

IBM

S#	Date	Speaker	Channel	URL
S1	2022-04-05	Ric Lewis, Senior Vice President, IBM Systems	Corporate press release, IBM Newsroom — launch of IBM z16 mainframe with on-chip Telum AI accelerator	https://newsroom.ibm.com/2022-04-05-Announcing-IBM-z16-Real-time-AI-for-Transaction-Processing-at-Scale-and-Industrys-First-Quantum-Safe-System
S2	2022-05-19	Tom Rosamilia, Senior Vice President, IBM Software	Corporate press release, IBM Newsroom — IBM Global AI Adoption Index 2022	https://newsroom.ibm.com/2022-05-19-Global-Data-from-IBM-Shows-Steady-AI-Adoption-as-Organizations-Look-to-Address-Skills-Shortages,-Automate-Processes-and-Encourage-Sustainable-Operations
S3	2023-05-01	Arvind Krishna, Chairman and CEO, IBM	Bloomberg News interview; exact words reproduced verbatim by Axios (published 2023-05-02)	https://www.axios.com/2023/05/02/ibm-stock-artificial-intelligence-hiring
S4	2023-05-09	Arvind Krishna, Chairman	Corporate press release, IBM Newsroom	https://newsroom.ibm.com/2023-05-09-IBM-Unveils-the-Watsonx-Platform-to-Power-Next-Generation-Foundation-Models-for-Business

		and CEO, IBM	— watsonx platform launch at THINK 2023	
S5	2023- 08-22	Arvind Krishna, Chairman and CEO, IBM	CNBC interview; exact words reproduced verbatim by Fortune (published 2023-08-22)	https://fortune.com/2023/08/22/ibm-ceo-arvind-krishna-ai-shrinking-workforces/
S6	2023- 10-25	Arvind Krishna, Chairman and CEO, IBM	IBM Q3 2023 earnings call, prepared remarks (The Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2023/10/25/international-business-machines-ibm-q3-2023-earnin/
S7	2024- 01-24	Arvind Krishna, Chairman and CEO, IBM	IBM Q4 2023 earnings call, prepared remarks (The Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2024/01/24/international-business-machines-ibm-q4-2023-earnin/
S8	2024- 01-24	Arvind Krishna, Chairman and CEO, IBM	IBM Q4 2023 earnings call, Q&A (The Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2024/01/24/international-business-machines-ibm-q4-2023-earnin/
S9	2024- 03-11	Arvind Krishna, Chairman and CEO, IBM	Chairman's letter, IBM 2023 Annual Report (Form ARS, filed with SEC 2024-03-11)	https://www.sec.gov/Archives/edgar/data/0000051143/000110465924032653/tm23326704d5_arss.pdf
S10	2024- 05-21	Arvind Krishna, Chairman and CEO, IBM	Corporate press release, IBM Newsroom — THINK 2024, open- sourcing of Granite models	https://newsroom.ibm.com/2024-05-21-IBM-Unveils-Next-Chapter-of-watsonx-with-Open-Source,-Product-Ecosystem-Innovations-to-Drive-Enterprise-AI-at-Scale
S11	2025- 05-05	Arvind Krishna, Chairman and CEO, IBM	Wall Street Journal interview; exact words reproduced verbatim by Entrepreneur (published 2025-05-08)	https://www.entrepreneur.com/business-news/ibm-ceo-ai-replaced-hundreds-of-human-resources-staff/491341
S12	2025- 05-06	Arvind Krishna, Chairman and CEO, IBM	Corporate press release, IBM Newsroom — THINK 2025, AI agent capabilities in watsonx Orchestrate	https://newsroom.ibm.com/2025-05-06-ibm-accelerates-enterprise-gen-ai-revolution-with-hybrid-capabilities
S13	2025- 10-20	Rob Thomas,	Corporate press release,	https://newsroom.ibm.com/2025-10-20-ibm-and-groq-partner-to-accelerate-enterprise-ai-deployment-with-speed-and-scale

		Senior Vice President, Software and Chief Commercial Officer, IBM	IBM Newsroom — IBM and Groq partnership on AI inference	
S14	2026-01-28	Arvind Krishna, Chairman and CEO, IBM	IBM Q4 2025 earnings call, prepared remarks (The Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2026/01/28/ibm-ibm-q4-2025-earnings-call-transcript/
S15	2026-01-28	Jim Kavanaugh, Senior Vice President and Chief Financial Officer, IBM	IBM Q4 2025 earnings call, prepared remarks (The Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2026/01/28/ibm-ibm-q4-2025-earnings-call-transcript/
S16	2026-04-22	Arvind Krishna, Chairman and CEO, IBM	IBM Q1 2026 earnings call, prepared remarks (The Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2026/04/22/ibm-ibm-q1-2026-earnings-call-transcript/

Salesforce

S#	Date	Speaker	Channel	URL
S1	2022-03-01	Gavin Patterson, President and Chief Revenue Officer	Q4 FY2022 earnings call (fiscal 2022 fourth quarter and full year results), prepared remarks; The Motley Fool transcript	https://www.fool.com/earnings/call-transcripts/2022/03/02/salesforcecom-crm-q4-2022-earnings-call-transcript/
S2	2022-09-20	Corporate press release	Salesforce newsroom press release announcing Salesforce Genie at Dreamforce 2022	https://www.salesforce.com/news/press-releases/2022/09/20/genie-news/
S3	2022-09-20	David Schmaier, President and Chief Product Officer	Salesforce newsroom press release announcing Salesforce Genie at Dreamforce 2022	https://www.salesforce.com/news/press-releases/2022/09/20/genie-news/
S4	2023-03-07	Marc Benioff, CEO of Salesforce	Corporate press release launching Einstein GPT, described as the first generative AI for CRM	https://www.salesforce.com/news/press-releases/2023/03/07/einstein-generative-ai/
S5	2023-08-30	Marc Benioff, Chairman and Co-Chief Executive Officer	Q2 FY2024 earnings call, prepared remarks; The Motley Fool transcript	https://www.fool.com/earnings/call-transcripts/2023/08/30/salesforce-crm-q2-2024-earnings-call-transcript/
S6	2023-08-30	Brian Millham, President and Chief Operating Officer	Q2 FY2024 earnings call, prepared remarks on customer deployments; The Motley Fool transcript	https://www.fool.com/earnings/call-transcripts/2023/08/30/salesforce-crm-q2-2024-earnings-call-transcript/

S7	2023-09-12	Marc Benioff, CEO of Salesforce	Dreamforce 2023 keynote, San Francisco; direct quotes reported by ITPro (article published 2023-09-13)	https://www.itpro.com/technology/artificial-intelligence/dreamforce-2023-benioff-says-generative-ai-needs-a-trust-revolution
S8	2024-09-12	Marc Benioff, Chair and CEO, Salesforce	Corporate press release unveiling Agentforce autonomous AI agents ahead of Dreamforce 2024	https://www.salesforce.com/news/press-releases/2024/09/12/agentforce-announcement/
S9	2024-12-03	Marc Benioff, Co-Founder, Chairman, and Chief Executive Officer	Q3 FY2025 earnings call, prepared remarks; The Motley Fool transcript	https://www.fool.com/earnings/call-transcripts/2024/12/03/salesforce-crm-q3-2025-earnings-call-transcript/
S10	2024-12-03	Brian Millham, President and Chief Operating Officer	Q3 FY2025 earnings call, prepared remarks; The Motley Fool transcript	https://www.fool.com/earnings/call-transcripts/2024/12/03/salesforce-crm-q3-2025-earnings-call-transcript/
S11	2024-12-03	Amy Weaver, President and Chief Financial Officer	Q3 FY2025 earnings call, prepared remarks; The Motley Fool transcript	https://www.fool.com/earnings/call-transcripts/2024/12/03/salesforce-crm-q3-2025-earnings-call-transcript/
S12	2025-02-26	Marc Benioff, Co-Founder, Chairman, and Chief Executive Officer	Q4 FY2025 earnings call, Q&A response on labor augmentation and replacement; The Motley Fool transcript	https://www.fool.com/earnings/call-transcripts/2025/02/26/salesforce-crm-q4-2025-earnings-call-transcript/
S13	2025-02-26	Marc Benioff, Co-Founder, Chairman, and Chief Executive Officer	Q4 FY2025 earnings call, Q&A response on labor augmentation; The Motley Fool transcript	https://www.fool.com/earnings/call-transcripts/2025/02/26/salesforce-crm-q4-2025-earnings-call-transcript/
S14	2025-07-30	Marc Benioff, CEO of Salesforce	Interview with Fortune (published 2025-07-30), quoted in Fortune article of 2025-09-02 syndicated on Yahoo Finance (fetched page)	https://finance.yahoo.com/news/salesforce-ceo-marc-benioff-says-145324020.html
S15	2025-09-02	Marc Benioff, CEO of Salesforce	Remarks on The Logan Bartlett Show podcast, reported with direct quotes by Fortune (2025-09-02), syndicated on Yahoo Finance (fetched page)	https://finance.yahoo.com/news/salesforce-ceo-marc-benioff-says-145324020.html
S16	2026-02-25	Marc Benioff, Chair and CEO, Salesforce	Q4 and full-year fiscal 2026 results press release (SEC Form 8-K, Exhibit 99.1)	https://www.sec.gov/Archives/edgar/data/0001108524/000110852426000056/crm-q4fy26xexhibit991.htm
S17	2026-04-29	Corporate newsroom story (Salesforce News, "How Salesforce Is Reshaping Its Workforce in the Age of AI")	Salesforce newsroom feature on internal Agentforce deployment and workforce redeployment	https://www.salesforce.com/news/stories/salesforce-reshaping-workforce-in-age-of-ai/

S18	2026-04-29	Sanjeev Balakrishnan, EVP of Customer Success at Salesforce	Salesforce newsroom feature on workforce reshaping in the age of AI	https://www.salesforce.com/news/stories/salesforce-reshaping-workforce-in-age-of-ai/
-----	------------	---	---	---

Shopify

S#	Date	Speaker	Channel	URL
S1	2022-07-27	Harley Finkelstein, President	Q2 2022 earnings call, prepared remarks (Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2022/07/27/shopify-shop-q2-2022-earnings-call-transcript/
S2	2022-07-27	Harley Finkelstein, President	Q2 2022 earnings call, prepared remarks (Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2022/07/27/shopify-shop-q2-2022-earnings-call-transcript/
S3	2022-11-15	Tobi Lütke, CEO	The Knowledge Project podcast Ep. 152 with Shane Parrish, transcript on fs.blog (episode publication date per Apple Podcasts: November 15, 2022)	https://fs.blog/knowledge-project-podcast/tobi-lutke-2/
S4	2022-11-15	Tobi Lütke, CEO	The Knowledge Project podcast Ep. 152 with Shane Parrish, transcript on fs.blog	https://fs.blog/knowledge-project-podcast/tobi-lutke-2/
S5	2023-05-04	Tobi Lütke, CEO	Company memo "Important team and business changes" announcing 20% workforce reduction and Shopify Logistics sale, published on Shopify newsroom	https://www.shopify.com/news/important-team-and-business-changes
S6	2023-05-04	Harley Finkelstein, President	Q1 2023 earnings call, prepared remarks, same day as the layoff memo (Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2023/05/04/shopify-shop-q1-2023-earnings-call-transcript/
S7	2023-07-26	Miqdad Jaffer, Head of Product for AI	Email interview quoted in TechCrunch coverage of the Editions Summer '23 Sidekick and Shopify Magic announcements	https://techcrunch.com/2023/07/26/shopify-sidekick-is-like-chatgpt-but-for-ecommerce-merchants/
S8	2023-08-02	Harley Finkelstein, President	Q2 2023 earnings call, prepared remarks (Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2023/08/02/shopify-shop-q2-2023-earnings-call-transcript/
S9	2024-02-13	Harley Finkelstein, President	Q4 2023 earnings call, prepared remarks (Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2024/02/13/shopify-shop-q4-2023-earnings-call-transcript/
S10	2024-02-13	Harley Finkelstein, President	Q4 2023 earnings call, prepared remarks (Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2024/02/13/shopify-shop-q4-2023-earnings-call-transcript/
S11	2024-08-28	Corporate press release (Shopify newsroom, "by Shopify")	Announcement "Shopify welcomes new CTO Mikhail Parakhin"	https://www.shopify.com/news/mikhail-parakhin-cto
S12	2025-04-07	Tobi Lütke, CEO	Internal memo to staff, posted publicly by Lütke on X (April 7, 2025); wording reproduced verbatim by BetaKit, which dates the memo itself March 20	https://betakit.com/shopify-ceo-tobi-lutke-tells-employees-to-prove-ai-cant-do-the-job-before-asking-for-resources/
S13	2025-04-07	Tobi Lütke, CEO	Same internal memo posted to X; wording reproduced verbatim by Entrepreneur (X origin noted)	https://www.entrepreneur.com/business-news/shopify-ceo-implements-ai-hiring-policy-for-employees/489718
S14	2025-12-10	Vanessa Lee, VP of Product	Shopify newsroom POV post announcing the Winter '26 Edition ("The Renaissance Edition")	https://www.shopify.com/news/winter-26-edition-renaissance
S15	2026-02-11	Harley Finkelstein, President	Q4 and full-year 2025 results press release, Shopify newsroom	https://www.shopify.com/news/shopify-q4-2025-financial-results
S16	2026-02-11	Jeff Hoffmeister, Chief Financial Officer	Q4 and full-year 2025 results press release, Shopify newsroom	https://www.shopify.com/news/shopify-q4-2025-financial-results

S17	2026-02-11	Harley Finkelstein, President	Q4 2025 earnings call, prepared remarks (Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2026/02/11/shopify-shop-q4-2025-earnings-call-transcript/
S18	2026-03-24	Mani Fazeli, VP Product	Shopify newsroom product announcement "Millions of merchants can sell in AI chats" (Agentic Storefronts / ChatGPT rollout)	https://www.shopify.com/news/agent-commerce-momentum

Duolingo

S#	Date	Speaker	Channel	URL
S1	2022-03-04	Corporate (Duolingo, Inc. — Annual Report, Form 10-K for FY2021)	SEC filing, Item 1 Business, "Application of data analytics and artificial intelligence to optimize learning"	https://www.sec.gov/Archives/edgar/data/1562088/000156208822000039/duol-20211231.htm
S2	2022-04-06	Corporate blog post (byline: Sophie Wodzak, Duolingo)	Duolingo blog, "Can a standardized test actually write itself?" on GPT-3 item generation for the Duolingo English Test	https://blog.duolingo.com/test-creation-machine-learning/
S3	2023-02-28	Luis von Ahn (co-founder and CEO)	Q4 2022 earnings call, answer on generative AI and content creation (Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2023/02/28/duolingo-duol-q4-2022-earnings-call-transcript/
S4	2023-02-28	Luis von Ahn (co-founder and CEO)	Q4 2022 earnings call, answer on human tutoring versus AI (Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2023/02/28/duolingo-duol-q4-2022-earnings-call-transcript/
S5	2023-03-14	Luis von Ahn (co-founder and CEO)	Corporate press release, "'Duolingo Max' Shows the Future of AI Education" (GPT-4 launch), Duolingo IR site	https://investors.duolingo.com/news-releases/news-release-details/duolingo-max-shows-future-ai-education
S6	2024-01-09	Duolingo company spokesperson (unnamed)	Statement to TechCrunch on the offboarding of roughly 10% of contractors at end of 2023	https://techcrunch.com/2024/01/09/duolingo-cut-10-of-its-contractor-workforce-as-the-company-embraces-ai/
S7	2024-02-28	Luis von Ahn (co-founder and CEO)	Q4 2023 earnings call, responding to press reports of AI-driven contractor cuts (Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2024/02/28/duolingo-duol-q4-2023-earnings-call-transcript/
S8	2025-04-28	Luis von Ahn (co-founder and CEO)	All-hands "AI-first" email posted to Duolingo's LinkedIn page; wording as reproduced verbatim by Mind the Product (LinkedIn origin)	https://www.mindtheproduct.com/duolingo-announces-ai-first-strategy/
S9	2025-04-30	Luis von Ahn (co-founder and CEO)	Corporate press release, "Duolingo Launches 148 New Language Courses" (GlobeNewswire)	https://www.globenewswire.com/news-release/2025/04/30/3071105/0/en/Duolingo-Launches-148-New-Language-Courses-Expands-Access-to-Popular-Languages-Including-Japanese-and-Korean.html
S10	2025-04-30	Jessie Becker (Senior Director of Learning Design, Duolingo)	Corporate press release, "Duolingo Launches 148 New Language Courses" (GlobeNewswire)	https://www.globenewswire.com/news-release/2025/04/30/3071105/0/en/Duolingo-Launches-148-New-Language-Courses-Expands-Access-to-Popular-Languages-Including-Japanese-and-Korean.html

S11	2025-05-23	Luis von Ahn (co-founder and CEO)	Follow-up LinkedIn post qualifying the AI-first memo; wording as reproduced verbatim by Entrepreneur (LinkedIn origin)	https://www.entrepreneur.com/business-news/duolingo-ceo-clarifies-ai-stance-after-backlash-read-memo/492141
S12	2025-08-06	Luis von Ahn (co-founder and CEO)	Q2 2025 earnings call, explaining DAU growth at the low end of guidance (Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2025/08/06/duolingo-duol-q2-2025-earnings-call-transcript/
S13	2026-02-26	Luis von Ahn (co-founder and CEO)	Q4 2025 earnings call, answer on becoming "AI native" (Motley Fool transcript; call held Thursday, Feb. 26, 2026)	https://www.fool.com/earnings/call-transcripts/2026/05/04/duolingo-duol-q4-2025-earnings-transcript/
S14	2026-04-13	Luis von Ahn (co-founder and CEO)	Remarks on the Silicon Valley Girl podcast, as quoted by Fortune, on dropping AI usage from performance reviews	https://fortune.com/2026/04/13/duolingo-ceo-luis-von-ahn-ai-usage-requirement-employee-performance-evaluations/
S15	2026-04-13	Duolingo company spokesperson (unnamed)	Statement to Fortune accompanying coverage of the performance-review reversal	https://fortune.com/2026/04/13/duolingo-ceo-luis-von-ahn-ai-usage-requirement-employee-performance-evaluations/

Cisco

S#	Date	Speaker	Channel	URL
S1	2022-05-04	Chuck Robbins, Chair and CEO	Corporate press release, "New Cisco Technology Can Predict Network Issues Before They Happen" (Cisco Newsroom), announcing an AI/ML predictive-analytics engine for networking	https://newsroom.cisco.com/c/r/newsroom/en/us/a/y2022/m05/new-cisco-technology-can-predict-network-issues-before-they-happen.html
S2	2022-08-17	Chuck Robbins, Chairman and CEO	Q4 FY2022 earnings call, prepared remarks (transcript published by The Motley Fool)	https://www.fool.com/earnings/call-transcripts/2022/08/17/cisco-systems-csco-q4-2022-earnings-call-transcript/
S3	2022-10-26	Corporate press release	WebexOne 2022 announcement, "Webex by Cisco Enhances Customer Experiences for the Digital Era" (Cisco Newsroom)	https://newsroom.cisco.com/c/r/newsroom/en/us/a/y2022/m10/webex-by-cisco-enhances-customer-experiences-for-the-digital-era.html
S4	2023-03-28	Jeetu Patel, EVP and GM, Security and Collaboration	Corporate press release, "Cisco Powers Unrivaled Hybrid Work	https://newsroom.cisco.com/c/r/newsroom/en/us/a/y2023/m03/cisco-powers-unrivaled-hybrid-work-experiences-with-purpose-built-ai-innovations-in-webex.html

			Experiences with Purpose-Built AI Innovations in Webex" (Cisco Newsroom)	
S5	2023-08-16	Chuck Robbins, Chairman and CEO	Q4 FY2023 earnings call, prepared remarks (transcript published by The Motley Fool, 2023-08-17)	https://www.fool.com/earnings/call-transcripts/2023/08/17/cisco-systems-csco-q4-2023-earnings-call-transcrip/
S6	2023-08-16	Chuck Robbins, Chairman and CEO	Q4 FY2023 earnings call, prepared remarks (transcript published by The Motley Fool, 2023-08-17)	https://www.fool.com/earnings/call-transcripts/2023/08/17/cisco-systems-csco-q4-2023-earnings-call-transcrip/
S7	2023-09-21	Chuck Robbins, Chair and CEO	Press release announcing the agreement to acquire Splunk for approximately \$28 billion (Exhibit 99.1 to Splunk Form 8-K, SEC)	https://www.sec.gov/Archives/edgar/data/1353283/000110465923102594/tm2326347d1_ex99-1.htm
S8	2024-04-18	Jeetu Patel, Executive Vice President and General Manager for Security and Collaboration	Corporate press release, "Cisco Reimagines Security for Data Centers and Clouds in Era of AI" — Hypershield launch (Cisco Newsroom)	https://newsroom.cisco.com/c/r/newsroom/en/us/a/y2024/m04/cisco-reimagines-security-for-data-centers-and-clouds-in-era-of-ai.html
S9	2024-08-14	Chuck Robbins, Chairman and CEO	Q4 FY2024 earnings call, prepared remarks, day the restructuring plan was announced (transcript published by The Motley Fool)	https://www.fool.com/earnings/call-transcripts/2024/08/14/cisco-systems-csco-q4-2024-earnings-call-transcrip/
S10	2024-08-14	Scott Herren, EVP and Chief Financial Officer	Q4 FY2024 earnings call, prepared remarks (transcript published by The Motley Fool)	https://www.fool.com/earnings/call-transcripts/2024/08/14/cisco-systems-csco-q4-2024-earnings-call-transcrip/
S11	2024-08-14	Scott Herren, EVP and Chief Financial Officer	Q4 FY2024 earnings call, Q&A response on whether layoff savings were embedded in guidance	https://www.fool.com/earnings/call-transcripts/2024/08/14/cisco-systems-csco-q4-2024-earnings-call-transcrip/

			(transcript published by The Motley Fool)	
S12	2024-08-14	Chuck Robbins, Chairman and CEO	Q4 FY2024 earnings call, Q&A, adding to Herren's answer on the restructuring (transcript published by The Motley Fool)	https://www.fool.com/earnings/call-transcripts/2024/08/14/cisco-systems-csco-q4-2024-earnings-call-transcrip/
S13	2025-01-15	Jeetu Patel, Executive Vice President and Chief Product Officer	Corporate press release, "Cisco Unveils AI Defense to Secure the AI Transformation of Enterprises" (Cisco Newsroom)	https://newsroom.cisco.com/c/r/newsroom/en/us/a/y2025/m01/cisco-unveils-ai-defense-to-secure-the-ai-transformation-of-enterprises.html
S14	2025-05-14	Chuck Robbins, chair and CEO	Q3 FY2025 earnings press release, "Cisco Reports Third Quarter Earnings" (Cisco Newsroom)	https://newsroom.cisco.com/c/r/newsroom/en/us/a/y2025/m05/cisco-reports-third-quarter-earnings.html
S15	2025-06-10	Jeetu Patel, President and Chief Product Officer	Corporate press release at Cisco Live San Diego, "Cisco Transforms Security for the Agentic AI Era, Further Fusing Security into the Network" (Cisco Newsroom)	https://newsroom.cisco.com/c/r/newsroom/en/us/a/y2025/m06/cisco-transforms-security-for-the-agentic-ai-era-further-fusing-security-into-the-network.html
S16	2026-02-10	Jeetu Patel, Cisco's President and Chief Product Officer	Corporate press release at Cisco Live EMEA, Amsterdam, "Cisco Redefines Security for the Agentic Era with AI Defense Expansion and AI-Aware SASE" (Cisco Newsroom)	https://newsroom.cisco.com/c/r/newsroom/en/us/a/y2026/m02/cisco-redefines-security-for-the-agentic-era.html
S17	2026-02-11	Charles (Chuck) Robbins, Chair and CEO	Q2 FY2026 earnings call, prepared remarks (transcript published by The Motley Fool)	https://www.fool.com/earnings/call-transcripts/2026/02/11/cisco-csco-q2-2026-earnings-call-transcript/
S18	2026-02-11	Charles (Chuck) Robbins, Chair and CEO	Q2 FY2026 earnings call, prepared remarks on internal AI adoption (transcript	https://www.fool.com/earnings/call-transcripts/2026/02/11/cisco-csco-q2-2026-earnings-call-transcript/

			published by The Motley Fool)	
--	--	--	-------------------------------------	--

Meta

S#	Date	Speaker	Channel	URL
S1	2022-01-24	Corporate blog post (AI at Meta blog; authored by Kevin Lee, Technical Program Manager, and Shubho Sengupta, Software Engineer)	Meta AI blog announcement of the AI Research SuperCluster (RSC)	https://ai.meta.com/blog/ai-rsc/
S2	2022-07-27	Mark Zuckerberg, Chief Executive Officer	Q2 2022 earnings call, prepared remarks (Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2022/07/27/meta-platforms-inc-meta-q2-2022-earnings-call-tran/
S3	2022-11-09	Mark Zuckerberg, Founder and CEO	Layoff letter to employees announcing ~11,000 job cuts, published on Meta Newsroom	https://about.fb.com/news/2022/11/mark-zuckerberg-layoff-message-to-employees/
S4	2023-02-24	Corporate blog post (AI at Meta blog, unbylined)	Meta AI blog announcing the first LLaMA model release	https://ai.meta.com/blog/large-language-model-llama-meta-ai/
S5	2023-03-14	Mark Zuckerberg, Founder and CEO	"Update on Meta's Year of Efficiency" letter to employees announcing ~10,000 further job cuts, Meta Newsroom	https://about.fb.com/news/2023/03/mark-zuckerberg-meta-year-of-efficiency/
S6	2023-07-18	Corporate press release	Meta Newsroom announcement of Llama 2 with Microsoft partnership	https://about.fb.com/news/2023/07/llama-2/
S7	2024-02-01	Mark Zuckerberg, CEO	Q4 2023 earnings call, prepared remarks (official transcript on Meta investor relations)	https://s21.q4cdn.com/399680738/files/doc_financials/2023/q4/META-Q4-2023-Earnings-Call-Transcript.pdf
S8	2024-04-18	Corporate press release	Meta Newsroom announcement of Meta AI assistant built on Llama 3	https://about.fb.com/news/2024/04/meta-ai-assistant-built-with-llama-3/
S9	2024-07-23	Mark Zuckerberg, Founder and CEO	"Open Source AI is the Path Forward" open letter published on Meta Newsroom, accompanying Llama 3.1 405B release	https://about.fb.com/news/2024/07/open-source-ai-is-the-path-forward/
S10	2025-01-20	Mark Zuckerberg, Founder and CEO	Remark on The Joe Rogan Experience podcast (January 2025), quoted verbatim in Entrepreneur article dated 20 January 2025 (podcast origin; fetchable reproduction logged)	https://www.entrepreneur.com/business-news/meta-developing-ai-engineers-to-write-code-instead-of-humans/485806
S11	2025-01-29	Mark Zuckerberg, Founder, Chair, and Chief Executive Officer	Q4 2024 earnings call, prepared remarks (Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2025/01/29/meta-platforms-meta-q4-2024-earnings-call-transcri/
S12	2025-01-29	Mark Zuckerberg, Founder, Chair, and Chief Executive Officer	Q4 2024 earnings call, prepared remarks (Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2025/01/29/meta-platforms-meta-q4-2024-earnings-call-transcri/

S13	2025-01-29	Susan Li, Chief Financial Officer	Q4 2024 earnings call, CFO outlook remarks (Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2025/01/29/meta-platforms-meta-q4-2024-earnings-call-transcri/
S14	2025-07-30	Mark Zuckerberg, Founder and CEO	"Personal Superintelligence" letter published at meta.com (linked from Meta Newsroom)	https://www.meta.com/superintelligence/
S15	2026-01-28	Mark Zuckerberg, Chief Executive Officer	Q4 2025 earnings call, prepared remarks (Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2026/01/28/meta-meta-q4-2025-earnings-call-transcript/
S16	2026-04-29	Mark Zuckerberg, Chief Executive Officer	Q1 2026 earnings call, prepared remarks (Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2026/04/29/meta-meta-q1-2026-earnings-call-transcript/
S17	2026-04-29	Mark Zuckerberg, Chief Executive Officer	Q1 2026 earnings call, prepared remarks (Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2026/04/29/meta-meta-q1-2026-earnings-call-transcript/
S18	2026-04-29	Susan Li, Chief Financial Officer	Q1 2026 earnings call, CFO segment remarks (Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2026/04/29/meta-meta-q1-2026-earnings-call-transcript/

Amazon

S#	Date	Speaker	Channel	URL
S1	2022-04-14	Andy Jassy, President and CEO	2021 Letter to Shareholders, published on aboutamazon.com (pre-ChatGPT baseline; Prime Video section)	https://www.aboutamazon.com/news/company-news/2021-letter-to-shareholders
S2	2022-06-23	Jeff Barr, Chief Evangelist, AWS	AWS News Blog post announcing the Amazon CodeWhisperer preview	https://aws.amazon.com/blogs/aws/now-in-preview-amazon-codewhisperer-ml-powered-coding-companion/
S3	2022-11-10	Corporate press release (bylined "Amazon Staff")	aboutamazon.com Operations announcement introducing the Sparrow warehouse robot	https://www.aboutamazon.com/news/operations/amazon-introduces-sparrow-a-state-of-the-art-robot-that-handles-millions-of-diverse-products
S4	2023-04-13	Andy Jassy, President and CEO	2022 Letter to Shareholders, published on aboutamazon.com	https://www.aboutamazon.com/news/company-news/amazon-ceo-andy-jassy-2022-letter-to-shareholders
S5	2023-04-27	Andy Jassy, President and CEO	Q1 2023 earnings call, opening remarks (The Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2023/04/28/amazoncom-amzn-q1-2023-earnings-call-transcript/
S6	2023-09-28	Swami Sivasubramanian, Vice President of Data and AI, AWS	Amazon newsroom announcement of Amazon Bedrock general availability (direct quote in release)	https://www.aboutamazon.com/news/aws/aws-amazon-bedrock-general-availability-generative-ai-innovations
S7	2024-02-01	Corporate blog post, bylined Rajiv Mehta (VP, Conversational Shopping) and Trishul Chilimbi (VP and Distinguished Scientist, Stores Foundational AI)	aboutamazon.com announcement launching Rufus in beta	https://www.aboutamazon.com/news/retail/amazon-rufus
S8	2024-04-11	Andy Jassy, President and CEO	2023 Letter to Shareholders, published on aboutamazon.com	https://www.aboutamazon.com/news/company-news/amazon-ceo-andy-jassy-2023-letter-to-shareholders
S9	2024-10-31	Andy Jassy, President and CEO	Q3 2024 earnings call, prepared remarks (The Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2024/10/31/amazoncom-amzn-q3-2024-earnings-call-transcript/
S10	2025-02-06	Andy Jassy, President and CEO	Q4 2024 earnings call, Q&A on capex (The Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2025/02/06/amazoncom-amzn-q4-2024-earnings-call-transcript/
S11	2025-04-10	Andy Jassy, President and CEO	2024 Letter to Shareholders, published on aboutamazon.com	https://www.aboutamazon.com/news/company-news/amazon-ceo-andy-jassy-2024-letter-to-shareholders

S12	2025-06-17	Andy Jassy, President and CEO	Message to employees, "Some thoughts on Generative AI," published in full on aboutamazon.com	https://www.aboutamazon.com/news/company-news/amazon-ceo-andy-jassy-on-generative-ai
S13	2025-10-28	Beth Galetti, Senior Vice President of People Experience and Technology	Message to employees announcing corporate workforce reduction of approximately 14,000 roles, published on aboutamazon.com	https://www.aboutamazon.com/news/company-news/amazon-workforce-reduction
S14	2026-02-05	Andy Jassy, President and CEO	Q4 2025 earnings call, prepared remarks (The Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2026/02/05/amazon-amzn-q4-2025-earnings-call-transcript/
S15	2026-04-09	Andy Jassy, President and CEO	2025 Letter to Shareholders, published on aboutamazon.com	https://www.aboutamazon.com/news/company-news/amazon-ceo-andy-jassy-2025-letter-to-shareholders
S16	2026-04-29	Andy Jassy, President and CEO	Q1 2026 earnings call, prepared remarks (The Motley Fool transcript)	https://www.fool.com/earnings/call-transcripts/2026/04/29/amazon-amzn-q1-2026-earnings-call-transcript/

JPMorgan Chase

S#	Date	Speaker	Channel	URL
S1	2022-04-04	Jamie Dimon, Chairman and CEO	Chairman & CEO letter to shareholders, 2021 Annual Report, jpmorganchase.com	https://reports.jpmorganchase.com/investor-relations/2021/ar-ceo-letters.htm
S2	2022-04-04	Jamie Dimon, Chairman and CEO	Chairman & CEO letter to shareholders, 2021 Annual Report, jpmorganchase.com	https://reports.jpmorganchase.com/investor-relations/2021/ar-ceo-letters.htm
S3	2023-04-04	Jamie Dimon, Chairman and CEO	Chairman & CEO letter to shareholders, 2022 Annual Report ("AI, Data and Our Journey to the Cloud" section), jpmorganchase.com	https://reports.jpmorganchase.com/investor-relations/2022/ar-ceo-letters.htm
S4	2023-04-04	Jamie Dimon, Chairman and CEO	Chairman & CEO letter to shareholders, 2022 Annual Report, jpmorganchase.com	https://reports.jpmorganchase.com/investor-relations/2022/ar-ceo-letters.htm
S5	2023-10-02	Jamie Dimon, Chairman and CEO	Broadcast interview on Bloomberg TV; exact words carried verbatim by CNBC Make It (published 2023-10-03)	https://www.cnbc.com/2023/10/03/jpmorgan-ceo-jamie-dimon-says-ai-could-bring-a-3-day-workweek.html
S6	2023-11-10	Lori Beer, Global Chief Information Officer	Remarks to Bloomberg News, reproduced verbatim by Banking Dive (published 2023-11-10)	https://www.bankingdive.com/news/jpmorgan-generative-ai-chatgpt-goldman-morgan-stanley-rbc-lori-beer/699427/
S7	2024-04-08	Jamie Dimon, Chairman and CEO	Chairman & CEO letter to shareholders, 2023 Annual Report ("The Critical Impact of Artificial Intelligence" section), jpmorganchase.com	https://www.jpmorganchase.com/ir/annual-report/2023/ar-ceo-letters
S8	2024-04-08	Jamie Dimon, Chairman and CEO	Chairman & CEO letter to shareholders, 2023 Annual Report, jpmorganchase.com	https://www.jpmorganchase.com/ir/annual-report/2023/ar-ceo-letters
S9	2024-04-08	Jamie Dimon, Chairman and CEO	Chairman & CEO letter to shareholders, 2023 Annual Report, jpmorganchase.com	https://www.jpmorganchase.com/ir/annual-report/2023/ar-ceo-letters
S10	2024-05-20	Mary Callahan Erdoes, CEO, Asset & Wealth Management	JPMorgan Chase Investor Day remarks (Bloomberg-reported), carried verbatim by PYMNTS (published 2024-05-20)	https://www.pymnts.com/news/banking/2024/jpmorgan-all-new-employees-will-receive-ai-training/

S11	2024-09-10	Daniel Pinto, President and Chief Operating Officer	Barclays Global Financial Services Conference remarks, as reported by CIO Dive (published 2024-09-11)	https://www.ciodive.com/news/JPMorgan-Chase-LLM-Suite-generative-ai-employee-tool/726772/
S12	2025-04-28	Teresa Heitsenrether, Chief Data and Analytics Officer	Direct interview quote in Harvard Business School case 9-325-066, "JPMorganChase: Leadership in the Age of GenAI" (rev. April 28, 2025; reviewed and approved before publication by a company designate), hosted by CTO Forum	https://www.ctoforum.org/wp-content/uploads/2025/07/JPMorganChase-Leadership-in-the-Age-of-GenAI-TMS.pdf
S13	2025-06-03	Corporate press release (JPMorganChase technology blog; no individual named for this passage)	Corporate tech-blog post, "LLM Suite named 2025 'Innovation of the Year' by American Banker," jpmorganchase.com	https://www.jpmorganchase.com/about/technology/blog/llmsuite-ab-award
S14	2025-10-07	Jamie Dimon, Chairman and CEO	Broadcast interview on Bloomberg TV at the JPMorgan Tech Stars Conference, London; exact words carried verbatim by Fortune (published 2025-10-08)	https://fortune.com/2025/10/08/jamie-dimon-head-in-the-sand-ai-will-take-jobs-tip-iceberg/
S15	2025-10-07	Jamie Dimon, Chairman and CEO	Broadcast interview on Bloomberg TV at the JPMorgan Tech Stars Conference, London; exact words carried verbatim by Fortune (published 2025-10-08)	https://fortune.com/2025/10/08/jamie-dimon-head-in-the-sand-ai-will-take-jobs-tip-iceberg/
S16	2026-02-24	Jamie Dimon, Chairman and CEO	Remarks to investors at a JPMorgan company update event; exact words carried verbatim by Fortune (published 2026-02-25)	https://fortune.com/2026/02/25/jamie-dimon-society-prepare-ai-job-displacement/
S17	2026-04-06	Jamie Dimon, Chairman and CEO	Chairman & CEO letter to shareholders, 2025 Annual Report ("AI, data and technology are key to the future" section), jpmorganchase.com	https://www.jpmorganchase.com/ir/annual-report/2025/ar-ceo-letters
S18	2026-04-06	Jamie Dimon, Chairman and CEO	Chairman & CEO letter to shareholders, 2025 Annual Report, jpmorganchase.com	https://www.jpmorganchase.com/ir/annual-report/2025/ar-ceo-letters