AISOLV RESEARCH · NOVEMBER 2026

The State of AI Deployment 2026

HEALTHCARE AND LIFE SCIENCES EDITION

How the world's largest companies actually rolled out AI, and what a $10M to $100M business can copy. 80 case studies, the eight lessons the leaders learned the hard way, and worksheets to plan your own.

J&J THERMO FISHER HEALTHSPAN SALESFORCE M3 INSURANCE + OPERATORS YOUR SIZE

What the leaders did, so you don't have to learn it the expensive way.

We went through how the 150 largest companies in the world rolled out AI, plus hundreds of mid-sized operators, and kept the 80 cases that hold up: their blogs, earnings calls, executive interviews, conference talks and case studies. More than 1,600 documented implementations. This report keeps the best of them, and turns them into a plan you can use.

WHAT'S INSIDE

- **Part 1. What the leaders learned.** Eight lessons, each with the companies that learned it.
- **Part 2. What they built.** Real implementations by function: sales, support, finance, operations, HR and IT, engineering, manufacturing, documents.
- **Part 3. Companies your size.** $10M to $100M operators and what they built with small teams.
- **Part 4. Your plan.** Worksheets to pick your first project, test it and roll it out.

HOW TO READ THE NUMBERS

- **Every result is company-reported.** It is what the company, its executive or its vendor said in public. We did not audit it.
- **Every case has a source.** The numbered list at the back links to each one.
- **Projections are marked.** Where a number is a target or an estimate, we say so.
- **Big companies are the benchmark, not the template.** Each lesson ends with what it means at your size.

> The pattern is the same at every size: pick the right work, keep a person where it counts, measure honestly, and make it easy to use.

AI in healthcare and life sciences: the cases to read first

Clinical documents, contracts, patient and customer questions, with a person kept where it counts.

HEALTHCARE COMPANY, ~130,000 STAFF

Johnson & Johnson

J&J uses AI to prepare documents for regulators, including clinical trial reports, which traditionally take hundreds of hours to assemble.

CIO Jim Swanson told Reuters a clinical trial report can take 700 to 900 hours the traditional way, and that this has gone "from 700 hours to about 15 minutes".

SOURCE [1]

LIFE SCIENCES COMPANY, GLOBAL

Thermo Fisher Scientific

Legal needed key terms pulled from a large set of contracts, a job estimated at three weeks for several paralegals. A small team used generative AI to extract the terms, improved accuracy with legal and technical experts, then joined the results with finance data.

Senior Solutions Architect Harry Look says the first analysis was delivered in 48 hours at about 70% accuracy, "turning what would have been a 3-week project with several paralegals" into something actionable within 2 days. Accuracy rose to about 87%, and joining 500+ contracts with finance data surfaced a $1.3 million opportunity.

SOURCE [2]

UK SUPPLEMENTS BRAND SELLING DIRECT (ORKLA-OWNED), ~GBP 50M TURNOVER, 50-SEAT CONTACT CENTRE

Healthspan

Healthspan's 'Product Professor' chat assistant answers customer questions about which vitamins and supplements suit them from the company's product knowledge base, and always offers a handoff to the contact centre at the end of a chat. Questions it cannot answer are passed to marketing as signals of products customers want.

Head of IT Rob King reports 90% of product-related customer queries in chat are now fully automated.

SOURCE [3]

UK SUPPLEMENTS BRAND SELLING DIRECT (ORKLA-OWNED), ~GBP 50M TURNOVER, 50-SEAT CONTACT CENTRE

Healthspan

Healthspan's contact-centre platform drafts AI replies to common customer emails such as order-status questions. Every draft goes to an agent, who proofreads, edits or declines it before it is sent; nothing goes out automatically.

Healthspan's Head of IT, Rob King, reports nearly 30% (28% on the slide) of all customer emails are now answered with AI-drafted replies, clearing backlogs faster with a more consistent tone, and says contact-centre headcount was not cut.

SOURCE [4]

SOFTWARE COMPANY, ~76,000+ STAFF, 1M+ CUSTOMERS

Salesforce

Salesforce put its own AI agent in front of customer support on its help portal. At first the team celebrated that only 1% of conversations were handed to a human, until they read the conversations and found frustrated customers who wanted a person while the agent kept trying.

Salesforce reports it changed the design so the handoff rate rose to about 5%; Chief Digital Officer Joe Inzerillo says the people who needed a human got one faster and customer satisfaction went up. The help agent was resolving 84% of queries autonomously with a 5% cut in case volume at the time.

SOURCE [5]

CLOUD AND RETAIL COMPANY, ~350,000 CORPORATE STAFF

Amazon (AWS)

In December 2025 AWS engineers let Amazon's own AI coding agent, Kiro, make changes using an account with broader permissions than expected. The agent reportedly decided to "delete and recreate the environment".

The Financial Times, via Engadget and The Register, reports a 13-hour outage of one AWS service, mainly affecting a region in China. Amazon called it "user error, not AI error", a "user access control issue", and said Kiro asks for authorization before acting by default.

SOURCE [6]

More from healthcare

INSURANCE AGENCY (WISCONSIN), ~$38M-$75M REVENUE, ~300-500 STAFF

M3 Insurance

M3's employee-benefits team uses an AI tool (Stitch Studio) to build the health plan comparison documents clients need before renewal, treating the AI like a junior colleague whose work is always checked. Carriers sometimes return proposals only about three days before the client meeting.

M3 reports a comparison that took 90-120 minutes to build now takes about 15 minutes; the freed time goes to client explanation and to seniors coaching junior staff.

SOURCE [7]

VENDOR-RISK SOFTWARE AND SERVICES FIRM, ~$36-42M REVENUE

Venminder (Ncontracts)

Venminder's analysts extract key terms (dates, auto-renewal, clauses) from client contracts. With Caylent it added Claude on AWS to search documents in plain language, read scans and pull the fields into its required formats, processing documents in parallel.

Caylent reports extraction accuracy rose from 60% to 86%+ (against a 75% target), about 70% of analysts' time was freed, and the contract review backlog fell from 65 days to 3 days within 4 days of adding generative AI.

SOURCE [8]

PART 1

What the leaders learned

Eight lessons from the largest AI rollouts in the world, each with the companies that learned it and what it means for a business your size.

1.1 Make it useful, then let people choose it

1.2 Forcing it, or banning it, backfires

1.3 Pick the work before the tool

1.4 Put a meter on it

1.5 Keep a person where it counts

1.6 Measure it properly

1.7 Fix the data first

1.8 Activity is not return

Make it useful, then let people choose it

The rollouts that stuck made AI easy to reach, trained people properly and tracked who used it. Few of them forced it.

BANK, ~317,000 STAFF

JPMorgan Chase

Built LLM Suite, an internal platform that gives staff access to outside AI models plus assistants they configure themselves. It was opt-in: the bank invited people rather than forcing it, and let peer stories and office-to-office competition do the selling.

JPMorgan's Derek Waldron reports going from zero to about 250,000 users in the first months, with one in two employees using it nearly every day. About 25-30% adopted straight away; growth to "a little north of 60%" came week by week after that.

> Launch AI as something people want to join, then expect the first quarter of staff to adopt on their own and the rest to need stories, training and nudges.

SOURCE [9]

BANK, ~120,000 STAFF

BBVA

Rolled out ChatGPT Enterprise and let employees build and share their own custom GPTs through an internal GPT store, backed by champions and hands-on training for 250 senior leaders. Every leader gets a monthly adoption dashboard showing their own, their peers' and their unit's usage, the CEO and chairman included.

OpenAI's BBVA customer story reports 20,000+ custom GPTs created (about 4,000 used frequently), 83% weekly active usage and about 3 hours saved per employee per week. BBVA's Antonio Bravo says more than 100 GPTs save about 80% of task time for 1,000+ employees.

> Make adoption visible to every manager, including yourself; a monthly usage dashboard creates healthy peer pressure without a mandate.

SOURCE [10]

BANK, ~213,000 STAFF

Bank of America

Launched Erica for Employees in 2020, an internal assistant for IT support such as password resets and device activation, then widened it in 2023 to HR questions such as benefits and payroll forms.

Bank of America reports over 90% of employees use Erica for Employees and it has cut calls into the IT service desk by more than 50%.

> Start with the dullest, highest-volume internal questions; a help-desk assistant everyone needs gets near-universal use and a clear saving.

SOURCE [11]

JOBS MARKETPLACE, ~20,000 STAFF IN RECRUIT'S HR TECHNOLOGY SEGMENT, 2,000+ ENGINEERS TRAINED

Indeed (Recruit Holdings)

Indeed first tried a champions model: about 60 'AI Coding Ambassadors' trained their teams, but teammates' use fell once the programme ended. It replaced this with AI Coding Essentials, a structured hands-on course that engineers were strongly encouraged (not mandated) to complete, with managers able to see completion.

Indeed reports weekly AI coding tool use rose from about 25% at the start of 2025 to about 97%. Engineers who completed the course cut coding time by about 35-36%; those who did not saw little change; coding time fell about 20% across the organisation.

> Access and champions are not enough; a short, structured course that managers track is what moves usage from a quarter of staff to nearly everyone.

SOURCE [12]

WHAT THIS MEANS AT $10M TO $100M

- ✓ Start with the 10-20 people who will actually use it, not a company-wide mandate.
- ✓ Train with real tasks from their week, not a generic course.
- ✓ Track weekly use by team, and talk about it with the leaders.

YOUR MOVE: _____

Forcing it, or banning it, backfires

Mandates without a good experience lost trust, and bans pushed people to use AI in secret. The fix was to redesign around the people using it.

TECH AND SERVICES COMPANY, ~280,000 STAFF

IBM

In 2018 IBM told more than 20,000 managers to stop going to HR and use its AskHR chatbot instead. Satisfaction collapsed, so IBM asked managers directly what they needed and rebuilt AskHR around that feedback, adding real transactions such as transfers and time off with a human tier for complex cases.

Fortune reports HR's score went from +19 before the mandate to an NPS of -35 within the first year. IBM now reports AskHR handles 94% of common questions without a human, with 11.5 million interactions in 2024 and 40% lower HR operating costs over four years.

> Forcing staff onto a tool built from the provider's point of view backfires; redesign around the users first, then adoption follows.

SOURCE [13]

ELECTRONICS MAKER (KOREA), GLOBAL WORKFORCE

Samsung Electronics

In May 2023 Samsung banned generative AI on company devices after engineers uploaded internal source code to ChatGPT. In 2026 it reversed course, giving ChatGPT Enterprise and Codex to all employees in Korea and all Device eXperience staff worldwide, with access tied to a mandatory internal AI security training and hands-on AI boot camps for about 2,300 executives.

No adoption or productivity numbers published. OpenAI calls it one of its largest enterprise deployments.

> A ban only pushes AI use underground; the safer path is an approved enterprise tool that people unlock by passing a short security course.

SOURCE [14]

INVESTMENT BANK, ~46,000 STAFF

Goldman Sachs

After letting traders and bankers use Claude Code and GitHub Copilot for months, Goldman updated its internal guidance in October 2026 to limit these coding tools to engineers and strats. Non-engineers were pointed to the firm's general AI assistant instead.

No numbers on how many people lost access. A Goldman spokesperson said that after experimentation weighing "productivity and risk management" the firm puts tools "in the hands of those who are best suited to leverage them for the long term."

> Giving everyone every tool is not the goal; match powerful tools to the roles that can use them safely, and give everyone else a simpler assistant.

SOURCE [15]

WHAT THIS MEANS AT $10M TO $100M

- ✓ If adoption is low, fix the experience before you push harder.
- ✓ Write a simple rule for what data can and cannot go into AI tools.
- ✓ Gate access with a short training, not a ban.

YOUR MOVE: _____

Pick the work before the tool

Not every task suits AI. The leaders cut hundreds of pilots down to a handful that mattered, and paired AI with plain automation where that was enough.

CYBERSECURITY COMPANY, ~20,000 STAFF

Palo Alto Networks

The CEO wanted generative AI to resolve up to 90% of the 480,000 employee IT, HR and finance requests a year. Analysis showed most requests asked for an action, not an answer, so the team built Panda, an agent that pairs AI answers with automation that actually does the task, and retired the phone line, Slack channel and portal.

Palo Alto Networks found AI alone, by summarising information, could solve only 18% of requests. CIO Meerah Rajavel reports close to 60% of employee tickets are now handled autonomously, and says this could rise to 80% over time (projection).

> Sort your requests by type before setting an AI target; most work needs a task done, not a question answered, so pair AI with automation.

SOURCE [16]

AIRCRAFT MAKER, ~157,000 STAFF

Airbus

Airbus trained 70,000-80,000 people on AI and launched hundreds of small generative AI pilots and proofs of concept. Judging most of them inefficient, it cut the portfolio down to the projects that could significantly change efficiency.

CEO Guillaume Faury says about 600 projects were launched and were "not very efficient", a "necessary pain" that reduced fear of AI, so Airbus reduced them to roughly 20 projects it now focuses on.

> Early pilots are cheap learning, but value comes from picking a handful of high-impact projects and stopping the rest.

SOURCE [17]

BANK, ~185,000 STAFF

Santander

After hundreds of pilots that mostly never reached production and 80+ separate chatbot efforts, Santander's AI team talked to every business line in every country, mapped the ideas and grouped them into three big bets: process automation, copilots for the sales force, and conversational agents for digital channels and the contact centre.

Chief AI Officer Inaki Bernal says more than 90% of the bank's early pilots never reached production, and that the roughly 1,000 AI initiatives identified were consolidated into three big bets.

> List every AI idea in the business, then group them into two or three bets you fund properly instead of running dozens of disconnected pilots.

SOURCE [18]

TOBACCO AND NICOTINE PRODUCTS COMPANY, GLOBAL

Philip Morris International

PMI opened bottom-up generative AI experimentation and ended up with 450 initiatives but no benefits at scale. It inventoried them and grouped them into 8-10 reusable capabilities (such as content generation, virtual personas and a first-pass legal and compliance check) built once on a central cloud platform and reused per use case.

PMI's speaker estimates the shared capabilities remove 70-80% of the effort for each new use case, which can also be delivered 40% faster than building from scratch.

> Most AI ideas share a few underlying building blocks; build those once and each new use case becomes a small job.

SOURCE [19]

WHAT THIS MEANS AT $10M TO $100M

- ✓ List every idea, then cut to the two or three with the most hours or money behind them.
- ✓ Check which steps really need judgement and which are plain rules.
- ✓ Build the shared pieces once and reuse them across projects.

YOUR MOVE: _____

Put a meter on it

AI bills grow quietly. The leaders set budgets per person, sent easy work to cheaper models and stopped projects that did not pay back.

CONSULTING FIRM, ~780,000 STAFF

Accenture

After centrally funding Claude Enterprise and Claude Code, Accenture reportedly moved to consumption-based funding from 1 October 2026: a set amount per person is funded centrally and heavier use is charged to project budgets.

404 Media, as recapped in news videos, reports a $29 per person per month central allowance after some Claude Code power users spent $500-1,000 a month in tokens. Accenture's spokesperson disputed the 'cap' framing, saying the figure "was used to demonstrate one part of the funding approach."

> Heavy AI users can cost 20-30 times a seat licence; decide early who pays for usage above the base allowance.

SOURCE [20]

TELECOM, 100,000+ STAFF ON ITS INTERNAL AI PLATFORM

AT&T

AT&T runs all AI use through an internal 'AI gateway' that routes simple jobs to cheaper models and tracks usage per person. It rebuilt its agent stack so large models hand narrow tasks to small, tuned models, and gives each user of coding agents like Claude Code or Codex a monthly 'wallet' that needs manager approval to exceed.

Chief Data Officer Andy Markus told VentureBeat the rebuilt orchestration produced up to 90% cost savings while token volume grew to 27 billion a day.

> Put one gateway in front of all AI use so you can see spend per person, route easy work to cheap models and make overspending a manager decision.

SOURCE [21]

NETWORKING COMPANY, ~90,000 STAFF

Cisco

Cisco gave its engineers unlimited AI tokens for the first four to six months of AI coding so they could get fluent. When costs rose sharply it built a 'tokenomics' app that shows what each agent is using, flags unusual spend and lets a manager pull a kill switch on a runaway agent, and gave every manager a token budget alongside the people budget.

No savings figures disclosed. President Jeetu Patel describes the rollout to the 32,000 people in his organisation as in progress (Oct 2026).

> Let people learn on generous limits first, then budget AI like headcount: each manager owns a token budget and can switch off agents that do not pay back.

SOURCE [22]

INDUSTRIAL GASES COMPANY, ~65,000 STAFF, ~EUR 27B REVENUE

Air Liquide

Air Liquide gave Gemini to all 65,000 employees and runs bigger AI projects under strict rules: 80% of AI resources go to the core business, projects are led by the business with an executive-committee-level sponsor, and the sponsor commits to the projected benefit before any large rollout.

Amaury de Menonville, Air Liquide's AI Acceleration Programme Director, says the sponsor's committed targets are tracked through standard group performance metrics, and the company stops about 20% to 30% of projects early.

> Make a named business owner sign up to the benefit before you scale an AI project, and plan to stop a quarter of them.

SOURCE [23]

WHAT THIS MEANS AT $10M TO $100M

- ✓ Give every tool and team a monthly budget from day one.
- ✓ Use the cheapest model that passes your tests, and step up only where it falls short.
- ✓ Ask each project owner to commit the benefit up front, and stop the ones that miss.

YOUR MOVE: _____

Keep a person where it counts

The safest rollouts let AI draft and a person decide. They set limits on what agents can touch and kept a way out to a human.

SOFTWARE COMPANY, ~76,000+ STAFF, 1M+ CUSTOMERS

Salesforce

Salesforce put its own AI agent in front of customer support on its help portal. At first the team celebrated that only 1% of conversations were handed to a human, until they read the conversations and found frustrated customers who wanted a person while the agent kept trying.

Salesforce reports it changed the design so the handoff rate rose to about 5%; Chief Digital Officer Joe Inzerillo says the people who needed a human got one faster and customer satisfaction went up. The help agent was resolving 84% of queries autonomously with a 5% cut in case volume at the time.

> A very low human-handoff rate can mean the AI is trapping customers; make reaching a person easy and judge success on satisfaction, not deflection.

SOURCE [5]

CLOUD AND RETAIL COMPANY, ~350,000 CORPORATE STAFF

Amazon (AWS)

In December 2025 AWS engineers let Amazon's own AI coding agent, Kiro, make changes using an account with broader permissions than expected. The agent reportedly decided to "delete and recreate the environment".

The Financial Times, via Engadget and The Register, reports a 13-hour outage of one AWS service, mainly affecting a region in China. Amazon called it "user error, not AI error", a "user access control issue", and said Kiro asks for authorization before acting by default.

> An AI agent can do anything its login allows; limit agent permissions and require a human sign-off before it changes anything live.

SOURCE [6]

BANK, ~180,000 STAFF WITH AI ACCESS, 30,000-40,000 DEVELOPERS

Citigroup

Citi uses an in-house tool, Citi Squad, for automated first-pass code reviews and gives developers the Devin coding agent for simple, repeatable jobs such as upgrades, patches, documentation and tests. The agent produces changes for a human; automated tests and a human review come before a developer accepts anything.

CEO Jane Fraser says AI has freed about 100,000 hours of developer capacity a week. CTO David Griffiths: "We're certainly not allowing any kind of deployment of code by the agents." He estimates Copilot users gain 5-15% and agents run 2x to 20x faster on suitable tasks.

> Let agents do the boring work and hand back a draft, but keep the final 'ship it' step with a person.

SOURCE [24]

UK SUPPLEMENTS BRAND SELLING DIRECT (ORKLA-OWNED), ~GBP 50M TURNOVER, 50-SEAT CONTACT CENTRE

Healthspan

Healthspan's contact-centre platform drafts AI replies to common customer emails such as order-status questions. Every draft goes to an agent, who proofreads, edits or declines it before it is sent; nothing goes out automatically.

Healthspan's Head of IT, Rob King, reports nearly 30% (28% on the slide) of all customer emails are now answered with AI-drafted replies, clearing backlogs faster with a more consistent tone, and says contact-centre headcount was not cut.

> Start AI in customer email as 'AI drafts, a person sends': you get speed and consistency without the risk of an unchecked reply.

SOURCE [4]

WHAT THIS MEANS AT $10M TO $100M

- ✓ AI drafts, a person sends: start every customer-facing or money-moving step this way.
- ✓ Give agents the smallest access that does the job, and log every action.
- ✓ Make it easy to reach a person. Customers notice when you do not.

YOUR MOVE: _____

Measure it properly

Usage is not value. The credible numbers came from control groups and honest baselines, and even big companies admit when the return is not there yet.

BANK, ~40,000 STAFF

DBS Group

Since 2021 DBS has put a dollar figure on its AI and data work each year. It measures value by comparing outcomes for customers who were offered AI-driven solutions against a control group who were not.

DBS reports about S$1 billion of economic value from AI, machine learning and data analytics in 2025, up from S$750 million in 2024, derived from comparisons between AI-enabled customers and control groups.

> If you want to prove AI value, hold back a control group and compare; otherwise you are counting activity, not results.

SOURCE [25]

SOCIAL MEDIA AND TECH COMPANY, ~78,000 STAFF

Meta Platforms

Meta pushes AI coding agents across engineering with adoption targets. Its CFO reports a headline productivity number on earnings calls, while its own developer-productivity team measured the effect of its DevMate agent more narrowly, comparing engineers by usage level.

CFO Susan Li said output per engineer rose 30% since the start of 2025, driven largely by AI coding agents (Q4 2025 earnings call). Meta's own measurement team found a 6-12% increase in code changes per developer per month for above-average users of DevMate, which it calls "the floor" for agentic coding on legacy code.

> Expect the careful internal measurement to come in far below the headline; ask how a number was measured before you plan around it.

SOURCE [26]

INDUSTRIAL TECHNOLOGY COMPANY, 105,000+ STAFF

ABB

ABB built Abby, an in-house assistant with a choice of 15+ models, and trained 24,000+ employees on AI in 2025. Its Global CIO is openly sceptical about measured returns so far.

ABB reports 35,000+ unique Abby users and 400,000+ conversations (2025). Global CIO Alec Joannou: "I'm going to challenge: no one's got a return on investment yet", adding that returns will come but need patience.

> Usage numbers are not returns; set expectations with your board that measurable payback takes time and has to be tracked deliberately.

SOURCE [27]

TECH AND SERVICES COMPANY, ~280,000 STAFF

IBM

IBM runs its 'client zero' programme, applying AI, automation and process removal across HR, IT, finance, procurement and supply chain. It measures productivity as the cost of shared services as a percentage of revenue.

CFO Jim Kavanaugh said IBM exited 2024 at $3.5 billion of annual run-rate savings and expected about $4.5 billion by end-2025 (projection at the time; IBM now says it 'unlocked' $4.5B). Next Platform notes IBM does not disclose its own AI training and inference costs, so net savings cannot be calculated.

> A gross savings figure without the AI bill against it is not ROI; track what the AI itself costs you.

SOURCE [28]

WHAT THIS MEANS AT $10M TO $100M

- ✓ Write down today's number before you build: hours, cycle time, error rate or cost.
- ✓ Compare a team using it with a team that is not, for a few weeks.
- ✓ Count the cost of running it, not only the time it saves.

YOUR MOVE: _____

Fix the data first

AI on messy data gives confident, wrong answers. Several leaders held AI back until their data and definitions were agreed.

OIL AND GAS MAJOR, $400B+ MARKET VALUE

ExxonMobil

ExxonMobil is moving a large, heavily customised SAP estate to standard processes and a 'clean core' on SAP S/4HANA, using process mining to map how work really happens, so that AI can be added later without a heavy lift. Before this, its data was fragmented and "every question was a science project".

No AI metric reported. The ExxonMobil executive on stage says the company is "doing some AI stuff" but is "happy to let that carry on over here on the side a little bit" until the fundamentals are "absolutely solid".

> If your core systems and data are a mess, AI will inherit the mess; clean up the process and data layer first and keep AI experiments small until then.

SOURCE [29]

AIRLINE, ~40 AIRCRAFT

Virgin Atlantic

Virgin Atlantic's BI team launched natural-language BI so leaders could ask questions and get answers in seconds. Engagement started high, but people then checked the answer on a dashboard or asked an analyst whether it was right. The fix was to write each metric's definition (such as NPS or on-time performance), get it signed off by a senior business owner and countersigned by the BI team, and make the AI answer only from those definitions.

No usage or accuracy figures disclosed. Head of BI Tom Barber calls the stall a "trust gap": "The governance was there, but the trust experience wasn't." The team now generates about 80% of the assets for a new analytics product with AI coding tools working from the signed-off definitions.

> Before you let AI answer business questions, agree in writing what each key number means and who owns it.

SOURCE [30]

POWER MANAGEMENT MANUFACTURER, $24.9B SALES

Eaton

When moving to a new contract management system, Eaton's analytics team found the system's built-in AI extracted contract details with under 40% accuracy. They built their own flow that scans, translates and uses GPT-4o to pull 30+ fields (supplier, dates, payment terms, currency, contract type) from each supplier contract into clean data for the new system.

Eaton's director of analytics reports 30,000+ supplier contracts converted with accuracy "exceeding 94%", compared with under 40% for the built-in AI; the proof of concept on 300 documents took 2 days. Legal has now asked for about 90,000 customer contracts.

> Turning messy documents into clean, structured data is a strong first AI project, and a focused team can beat a vendor's bundled AI feature.

SOURCE [31]

WHAT THIS MEANS AT $10M TO $100M

- ✓ Agree on what your key numbers mean before anyone asks AI about them.
- ✓ Clean the one dataset your first project depends on, not everything.
- ✓ Test the tool your software vendor bundles before you assume it is good enough.

YOUR MOVE: _____

Activity is not return

Thousands of agents and high adoption can still add up to little. Most of the value tends to come from a few use cases.

TECH INVESTOR AND OPERATOR, ~40,000 STAFF ACROSS ITS PORTFOLIO COMPANIES

Prosus

Prosus staff built over 60,000 AI agents on its internal Toqan platform in about 18 months. It then analysed which agents mattered, measuring shared agents by hours saved, revenue gained or cost removed, and asking business leaders the 'delete it tonight' question: what would happen to revenue or costs if this agent were deleted?

Prosus reports a power law: about 2% of the most active agents drive a disproportionate share of business impact. Of agents measured in hours saved, 82% save under 20 hours a month, 17% save 20-173 hours (up to about one full-time employee), and under 1% save thousands of hours.

> Most AI agents will be small personal helpers; find the few that matter and put your effort behind those.

SOURCE [32]

TELECOM OPERATOR, ~20,000 EMPLOYEES TOOK PART

SoftBank Corp.

At a June 2025 all-hands the CEO built an agent live and set 'summer homework': every employee should build 100 AI agents by the end of August. After the rush, the company found most agents could not be used in real work and switched from quantity to quality, filtering down to the valuable ones.

SoftBank reports more than 2.5 million agents created in about ten weeks, with 90% of survey respondents positive about the exercise. Business+IT reports most were unusable and that over 6,000 agents were kept after selection.

> Agent counts measure enthusiasm, not value; a build-a-thon is good for learning but plan the clean-up and selection step from day one.

SOURCE [33]

SOFTWARE AND CLOUD COMPANY, ~160,000 STAFF AT ROLLOUT

Oracle

Oracle set standards, security controls and policies, then rolled out ChatGPT Enterprise and Codex across the company. It now shows employees which model they are using and what it costs, steering routine work to cheaper models.

Global CIO Jae Evans said 80% of staff were using the tools within three months, but "when the bills came in, the company was surprised by the costs". Developers now produce in a week what took a team two to three quarters, yet delivery to customers did not speed up because testing, validation and release processes became the bottleneck.

> Fast adoption brings a fast bill and moves the bottleneck downstream; budget for usage and fix the steps after the AI before celebrating.

SOURCE [34]

WHAT THIS MEANS AT $10M TO $100M

- ✓ Do not celebrate seat counts or agent counts. Ask what changed in the numbers.
- ✓ Find the few uses that carry most of the value, and invest there.
- ✓ Retire tools and agents nobody uses.

YOUR MOVE: _____

PART 2

What they built, by function

Real implementations from the leaders and from companies your size, grouped by the part of the business they changed.

2.1 Sales and quoting

2.2 Customer support

2.3 Finance and accounting

2.4 Orders and operations

2.5 HR and IT help desks

2.6 Engineering and software

2.7 Manufacturing and supply chain

2.8 Documents, legal and compliance

Sales and quoting

Quotes, proposals and lead follow-up: drafted by AI, checked and sent by people.

INDUSTRIAL MACHINERY MAKER (INDIA), ~100 STAFF, NO DATA SCIENCE TEAM

Machinecraft

The owner fed decades of quotes, drawings, payment schedules and email threads into a set of AI agents ('Ira') that reads the inbox, drafts sales emails and quotes, prepares pre-call briefs and revives dead leads. Off-the-shelf models were used with no training or fine-tuning, and a fact-checking agent and a record of human corrections keep answers grounded.

Founder Rushabh Doshi reports an agency quoted about $230,000 to build it; they built it for around $30,000 and it runs on "a couple of thousand dollars a month". The golden rule: "Ira drafts, human sends." No pipeline or win-rate results disclosed.

SOURCE [35]

FOODSERVICE EQUIPMENT DISTRIBUTOR, ~$1.3B REVENUE, ~340 SALES ASSOCIATES

Edward Don & Company

Edward Don used an AI matching tool to read large customer RFPs and competitor quotes and suggest the right catalog item (plus alternatives) for each line, applying customer-specific rules. Sellers accept or correct each match, the corrections teach the system, and a human check is always kept before a quote goes out.

CIO Tim Walter reports that in the pilot a 1,600-line RFP that took 5-6 weeks by hand was matched in 20 minutes at 86% accuracy. The company targets a 400%+ increase in RFPs answered and $20-30M of a $50-60M missed opportunity in year one (targets, not results).

SOURCE [36]

RESEARCH SPECIMENS SUPPLIER, MID-MARKET, 25,000+ QUOTES A YEAR

BioIVT

BioIVT built 'Ivy', an AI agent that reads incoming quote requests, matches products against 15,000+ product codes and builds the quote, which a service rep reviews and sends. It started with simple catalog items and expanded to configurable products; every rep edit is captured and used to improve the agent.

Senior Director LaRae Bennett reports quote turnaround went "from days to minutes", reps have 43% more time for high-value work, and the agent is at a 94% accuracy rate.

SOURCE [37]

B2B ARM OF UBER, ONBOARDING TEAM OF 30 SPECIALISTS

Uber for Business

Tens of thousands of business leads a month started but did not finish self-serve sign-up, and 30 onboarding specialists could not follow up with them all. An AI agent now spots where each prospect stalled, sends a personalised email from an approved knowledge base, answers common questions and hands complex deals to a rep with a summary.

Salesforce reports that within two weeks of launch the agent's emails converted about 60% higher than the organic baseline, near 100% of leads now get follow-up, and Uber anticipates a seven-figure incremental revenue impact (projection).

SOURCE [38]

BANK (JAPAN), PART OF SMFG

SMBC (Sumitomo Mitsui Banking Corporation)

SMBC began rolling out an app built with Sakana AI in which several AI agents split the work of a client proposal for large corporate clients: one gathers and analyses data, one builds the hypothesis and storyline, one checks quality and facts. Bankers then focus on analysis, solution design and the client conversation.

Projection only: Nikkei reports large-corporate proposals that took 1-2 weeks are expected to take tens of minutes to a few hours. No measured results yet.

SOURCE [39]

YOUR NOTES

Which of these would fit your business, and who would own it?

Customer support

Answers from your own content, with a clear hand-off to a person.

D2C WALLETS AND ACCESSORIES BRAND, $100M+ REVENUE, ~100 STAFF

Ridge

Ridge moved its customer service onto Richpanel with an AI agent answering email and chat tickets (order status, returns, product questions) and escalating the rest to a smaller human team. Ridge does not use voice AI.

CEO Sean Frank says Ridge grew about 50% year over year while human-answered tickets hit about a five-year low, and the CX team is now 10 agents where it used to be 25. In a conversation with Richpanel's CEO he put his cost per ticket at "like a quarter" (about $0.25).

SOURCE [40]

TELECOM, ~75,000 US STAFF

T-Mobile US

T-Mobile added AI agents to customer care, including voice AI that answers the phone, and aims to fix problems before customers need to call. When it first launched its chatbot it blocked screenshots; the CEO overruled this and told the team to let customers critique it publicly.

Julianne Roberson, T-Mobile's director of AI engineering, says voice AI now answers about half of customer service calls, with over 200,000 AI-driven interactions a day.

SOURCE [41]

TELECOM, ~100,000 STAFF

Verizon

Verizon gave 28,000 customer care and store reps a Google Gemini-based 'Personal Research Assistant' that gives real-time, personalised answers to customer questions, so reps no longer put customers on hold to search across tools.

Verizon says the assistant gives "95 percent comprehensive answerability" for customer inquiries (CX Today). Shankar Arumugavelu (Verizon) reports hold time down 20% and an 80% reduction in reps switching between tools.

SOURCE [42]

RESALE MARKETPLACE APP (US ARM OF MERCARI, INC.), ~120,000 SUPPORT TICKETS A MONTH

Mercari US

One CX systems manager built a ~90-step Zapier workflow that classifies every ticket, pulls the matching policy and approved answers from a Google Sheet, and lets AI write the reply only within that fixed path. Customer data is scrubbed before it reaches the model, and anything low-confidence goes back to a human.

Zapier reports Mercari now sends a first reply on about 97% of its roughly 120,000 monthly tickets and resolves about 47,000 a month with no human response, freeing 3,000+ support hours a month.

SOURCE [43]

ONLINE RETAILER (AUSTRALIA AND NZ), ASX-LISTED, $100M-$500M REVENUE BAND

Kogan.com

Kogan's customer-care team, not engineers, built an AI chat agent on Salesforce Agentforce, starting with 'where is my order' questions and later adding warranty and membership queries. The same back end now also runs a bot for its Mighty Ape brand.

Kogan reports the agent's customer satisfaction is within plus or minus 5 percentage points of human agents, customers choosing the chatbot rose from 58% to 78%, and the second brand's bot launched in under two weeks.

SOURCE [44]

YOUR NOTES

Which of these would fit your business, and who would own it?

Finance and accounting

Invoices, collections, close and reporting.

SOFTWARE COMPANY, ~223,000 STAFF

Microsoft

Microsoft's finance and IT teams built a collections agent on Dynamics 365. It works from an AI-prioritised list of customers likely to pay late, summarises each account, writes the reminder email and helps answer customers' questions so they pay on time.

Microsoft finance leader Cory Hrncirik says it took the time to resolve late-paying customers "from 26 days... to something like five or six today" (his words: "I think"), saving thousands of hours and cost on work that had been outsourced.

SOURCE [45]

PRINTING AND PACKAGING COMPONENTS MAKER, ~225 STAFF, ~$37M REVENUE, 7 LEGAL ENTITIES

Pamarco Technologies

Pamarco runs bounded AI agents on Microsoft Dynamics 365 for month-end close tasks, coding freight and German supplier invoices, and building scheduled reports such as the sales director's board deck. The agents email the finance team a summary of how each invoice was coded.

CIO Glyn Wilson reports roughly 40% of financial close tasks are now automated, with potential to reach 50-60% by year-end (projection). Freight invoice coding runs at about a 95% success rate, so the team turned on auto-posting for non-PO invoices.

SOURCE [46]

ON-SITE FUEL DELIVERY COMPANY (US), ~260-270 STAFF, MULTIPLE ENTITIES

Diesel Direct

Diesel Direct's accounts payable team adopted Vic.ai to read and code invoices, match bills of lading to purchase orders, route approvals and run payments, posting straight into its Great Plains ERP.

Vic.ai reports 84% of invoices processed with no human touch, 99% accuracy in invoice coding by month six, 65% less invoice processing time and AP close shortened from 4 to 3 days. VP and Corporate Controller Andrew Turkish says the team is at about 2.5 FTE of efficiency against a 3 FTE business case.

SOURCE [47]

ENTERPRISE TECH COMPANY, FORTUNE 500, FINANCE TEAM OF 3,000+

Hewlett Packard Enterprise

HPE's finance team built 'CFO Insights' (Alfred) with Deloitte, an agentic system that assembles the data and commentary for financial reporting, including the weekly operations review that used to need a 100-slide Monday deck.

Deloitte reports CFO Insights is cutting HPE's financial reporting cycle time by about 40% and processing costs by at least 25%. HPE's finance COO estimates about 90% of the manual prep for the weekly review has been removed (Fortune).

SOURCE [48]

DELIVERY PLATFORM, 7M+ US ORDERS A DAY; AP TEAM OF 10-15

DoorDash

DoorDash's data teams built a governed platform over Snowflake on which non-technical accountants describe a process in plain language and build their own bots. The bots pull files from email, answer Jira tickets and send Slack prompts with an approve button; a human clicks to proceed.

Chief Accounting Officer Gordon Lee says the AP team, "not that big, 10-15 people", identified roughly 4,500 processes globally and built about 80 bots in two weeks. Time per step is measured before and after to compute ROI; no savings figures given.

SOURCE [49]

YOUR NOTES

Which of these would fit your business, and who would own it?

Orders and operations

Orders arriving by email, PDF and phone, turned into clean records.

FAMILY-OWNED PAPER AND PACKAGING DISTRIBUTOR (OHIO), ~256 STAFF, ~5,000 CUSTOMERS

Millcraft Paper Company

Millcraft built an in-house AI model that reads the PDF purchase orders customers email in, pulls out the PO number, items and quantities, and creates the order in its system with no re-keying. Customer service staff moved to proactive calls, problem prevention and add-on sales; no jobs were cut.

President Travis Mlakar says about 60% of inbound orders now arrive electronically: about 20% via e-commerce, some by EDI, and about 30% of all transactions through the AI order-entry model. He says order mistakes are down.

SOURCE [50]

FREIGHT LOGISTICS COMPANY, 37M SHIPMENTS A YEAR, 83,000 CUSTOMERS

C.H. Robinson

C.H. Robinson built AI agents that read emailed shipping requests (often with attachments and handwritten notes), connect information across the email, fetch missing details and create the order. Previously an email could wait up to four hours in a queue and took about seven minutes to process by hand.

LangChain reports about 5,500 orders a day are now automated, saving over 600 hours a day on this task.

SOURCE [51]

ORDER SOFTWARE VENDOR SERVING 21,000+ FOOD DISTRIBUTORS

Choco (for food distributors)

Choco's AI order agent turns unstructured orders from restaurants (emails, PDFs, texts, voicemails) into orders in the distributor's ERP, with configurable thresholds for what goes straight through and what a person checks, and keeps learning from corrections.

OpenAI reports Choco processes over 8.8 million orders a year, with up to 50% less manual order entry and 2x productivity for distributors "without additional headcount", and error rates "below 1-5%".

SOURCE [52]

CONSUMER WIPES BRAND, ~$250M REVENUE RUN-RATE, 40 STAFF

Dude Wipes

One non-technical operations analyst, with an outside automation partner, built 100+ automations on Make. They include an AI agent ('Diesel Dan') that plans and creates all inbound shipments from manufacturers to warehouses, and an agent that reads emails and creates sample orders with tracking updates.

Make reports Dude Wipes is on track for $250M in annual revenue with 40 employees (from $72M with 14 people three years earlier). The shipment agent frees about 200 hours a year per supply chain team member; the sample-order agent gave one marketer back about 600 hours a year.

SOURCE [53]

PRESS-ON NAILS AND LASHES BRAND, ~$50M+ REVENUE, ~50-70 STAFF

Glamnetic

Glamnetic's support AI agent on Gorgias can act, not just answer: it edits shipping addresses, cancels orders and changes items before the warehouse ships them. Order changes were 11% of post-purchase tickets.

Gorgias reports 41% of order-modification requests are now resolved without a human (up from 19% on the older model), 2.2x more addresses edited in time and 1.8x more orders cancelled before fulfilment, and 44% fewer of these conversations needing a human, over three months.

SOURCE [54]

YOUR NOTES

Which of these would fit your business, and who would own it?

HR and IT help desks

The questions staff ask every day, answered without a ticket queue.

TECH AND SERVICES COMPANY, ~280,000 STAFF

IBM

IBM built AskIT in about 100 days as the first stop for every employee IT request, starting with a 12,000-user pilot trained on the issues behind 80% of tickets. Content specialists rewrote support articles so the AI could use them, and human agents now handle the complex cases.

IBM reports AskIT now resolves 86% of issues without human intervention, with CSAT above 90 (up 11.6 points since launch). IT support calls and chats fell from more than a million a year to fewer than 100,000 between 2021 and 2025.

SOURCE [55]

SEMICONDUCTOR AND SOFTWARE COMPANY, ~40,000 STAFF PLUS ~10,000 CONTRACTORS

Broadcom

Broadcom put a Moveworks AI assistant inside Google Chat as the front door for IT support, connected to its IT knowledge base, and used it to handle spikes in access and password requests when onboarding staff from acquisitions.

Moveworks reports autonomous resolution rose from about 10% at launch to 54%, and to 88% of support issues as of 2025, with more than 75,000 IT issues resolved and employees getting help in seconds instead of days.

SOURCE [56]

BANK, ~48,900 STAFF

Commonwealth Bank of Australia

CBA's ChatIT assistant in Microsoft Teams lets staff describe an IT problem in plain language, then runs targeted automations to fix it on their device, with no phone call or ticket.

Microsoft reports that in the six months after the automation function launched, ChatIT processed over 2.3 million messages and ran more than 12,000 automated fixes, each averaging about 2 minutes against a typical 17-minute service-desk call, saving nearly 2,500 hours.

SOURCE [57]

ENERGY MANAGEMENT AND AUTOMATION COMPANY, ~140,000-160,000 STAFF

Schneider Electric

Schneider built 'Jo', a chatbot in Microsoft Teams that answers employee HR and IT questions from approved documents. Because tolerance for wrong answers on HR and labour law is zero, the team tuned it to say it cannot help yet rather than guess.

No figures disclosed. Anna Gawlicka (Schneider Electric) says Jo now gives more replies telling employees it "is not yet able to support the question", which they see as "better user experience than providing an incorrect answer." Tracked metrics are users, meaningful conversations, answer accuracy and potentially saved tickets.

SOURCE [58]

Engineering and software

Code review, migrations and old systems nobody can read anymore.

RIDE-HAILING AND DELIVERY PLATFORM, ~65,000 CODE CHANGES A WEEK

Uber

Uber built uReview, an AI second reviewer that comments on every code change before a human reviewer sees it. It filters out low-confidence and low-value suggestions so engineers only get comments likely to matter.

Uber reports uReview analyses over 90% of its ~65,000 weekly code changes; engineers who rate it mark 75% of comments useful, and 65% of its comments are acted on in the same change, versus 51% for human-written review comments. It estimates about 1,500 hours saved a week, assuming 10 minutes of second-reviewer time per change.

SOURCE [59]

PAYMENTS COMPANY

Stripe

Stripe built 'Minions', unattended coding agents that take a task and produce a complete pull request in one shot, running in isolated developer environments. Every change is reviewed by a human before it merges, but contains no human-written code.

Stripe reports over 1,300 pull requests a week merged that are completely minion-produced and human-reviewed (up from 1,000 a few weeks earlier).

SOURCE [60]

AIRCRAFT MAKER, ~157,000 STAFF

Airbus

Airbus had 40 mainframe applications written in COBOL that, apart from the CIO herself, "absolutely nobody" could code in. It used AI to rebuild the documentation of what each application did and how its logic worked.

CIO Catherine Jestin says AI let them "completely rebuild the documentation" of the functional logic and algorithms, "which helped us to migrate to the next generation." No time or cost figures given.

SOURCE [61]

MUSIC STREAMING COMPANY

Spotify

Spotify added a background coding agent, Honk, to its Fleet Management platform, which applies the same code change across thousands of repositories. The agent handles migrations that simple scripts could not, such as code changes that need judgment, with automated checks before merge.

Spotify reports 1,500+ AI-generated pull requests merged into production and a total time saving of 60-90% on those migrations "compared to writing the code by hand".

SOURCE [62]

CREATOR E-COMMERCE PLATFORM, ~$12-24M REVENUE, ~110 STAFF

Gumroad

Gumroad's CEO prototypes a feature in v0, hands a short spec or screenshot to the Devin coding agent (often several variants at once), and finishes in Cursor if Devin falls short; each change gets a preview build and human review. Designers and PMs are encouraged to ship directly.

In the interview the host notes Devin is writing 41% of Gumroad's PRs with a target of 80% by year-end, which CEO Sahil Lavingia discusses as current. In one demo a change estimated at 24 hours merged in about an hour.

SOURCE [63]

YOUR NOTES

Which of these would fit your business, and who would own it?

Manufacturing and supply chain

Inspection, planning and the factory floor.

CONSUMER GOODS COMPANY, ~$84B REVENUE

Procter & Gamble

In Brazil, retailers place a month's order up front and ask P&G to deliver when ready. P&G had been sequencing deliveries by its own needs (excess stock, production runs) rather than demand, causing one of its worst out-of-stock problems. It used AI on supply planning, plant capacity, weather, seasonality and holiday data to decide which product goes to which store when, and fed that back into the production schedule.

CIO Seth Cohen says P&G improved out-of-stocks in Brazil by 15 percentage points, where normal improvements are 1-2 points.

SOURCE [64]

AUTOMAKER, GLOBAL

Toyota

Toyota built a factory AI platform so that line workers, regardless of AI skill, can create their own image-recognition and anomaly models for tasks like inspection, without data scientists.

Toyota reports the number of models created on the manufacturing floor rose from 8,000 in 2023 to 10,000 in 2024, saving over 10,000 man-hours a year in manufacturing.

SOURCE [65]

AIRCRAFT ENGINE MAKER, ~52,000 STAFF

GE Aerospace

GE Aerospace joined ERP, planning, supplier and engineering data into one model on Palantir so that parts likely to block an engine repair or delivery are flagged automatically, with suggested actions, instead of being found by hand.

A GE Aerospace leader says: "Every Monday, eight hours, their entire day was spent stitching together data into Excel spreadsheets to identify a handful of parts"; those leaders are now fed constraints in near real time.

SOURCE [66]

CONSUMER GOODS COMPANY, HUNDREDS OF FACTORIES

Unilever

Unilever builds AI-enabled digital twins of production lines, fed by live shop-floor data, that predict process problems and recommend settings to operators. It is scaling these to 40+ sites with Accenture.

Unilever reports its Raeford, North Carolina factory uses digital twins to achieve a 20% reduction in waste and a 10% uplift in capacity, and that at Hefei, China, AI contributed an 8% OEE increase and 20% less waste.

SOURCE [67]

PLASTICS MANUFACTURER (MICHIGAN), ~1,100 STAFF, $100M-$500M REVENUE

Cascade Engineering

When Cascade launched the Ford F-150 front trunk ('frunk'), a line team member suggested training two collaborative robots to sweep each part for defects within the cycle time, capturing far more data points than manual inspection could.

CEO Christina Keller says the cobot inspection has now run for almost three years "and there's been zero quality defects".

SOURCE [68]

YOUR NOTES

Which of these would fit your business, and who would own it?

Documents, legal and compliance

Contracts, filings, reports and checks.

INVESTMENT BANK, ~80,000 STAFF

Morgan Stanley

Morgan Stanley's institutional securities group built an 'IPO builder' that drafts the S-1 registration document a company files to go public.

Jeff McMillan (Morgan Stanley) says an S-1 "takes on average about 65 hours to build a prospectus using the old tools. Using this tool, I can build it in 30 minutes. And by the way, it's better."

SOURCE [69]

HEALTHCARE COMPANY, ~130,000 STAFF

Johnson & Johnson

J&J uses AI to prepare documents for regulators, including clinical trial reports, which traditionally take hundreds of hours to assemble.

CIO Jim Swanson told Reuters a clinical trial report can take 700 to 900 hours the traditional way, and that this has gone "from 700 hours to about 15 minutes".

SOURCE [1]

LIFE SCIENCES COMPANY, GLOBAL

Thermo Fisher Scientific

Legal needed key terms pulled from a large set of contracts, a job estimated at three weeks for several paralegals. A small team used generative AI to extract the terms, improved accuracy with legal and technical experts, then joined the results with finance data.

Senior Solutions Architect Harry Look says the first analysis was delivered in 48 hours at about 70% accuracy, "turning what would have been a 3-week project with several paralegals" into something actionable within 2 days. Accuracy rose to about 87%, and joining 500+ contracts with finance data surfaced a $1.3 million opportunity.

SOURCE [2]

FOOD AND BEVERAGE COMPANY, OPERATING IN 185 COUNTRIES

Nestle

Nestle's procurement team applies AI to its supplier contracts to find inconsistencies between global contract terms and how they are applied in local markets.

A Nestle spokesperson says the analysis covers over 80% of total spend and has reviewed hundreds of thousands of contracts. CIO Chris Wright says the company's business cases show "the biggest value is in buying better, or reducing energy in a factory, not only being more humanly efficient."

SOURCE [70]

BANK, ~211,000 STAFF, ~41M CUSTOMERS

HSBC

HSBC co-developed with Google a machine-learning system ('Dynamic Risk Assessment') that scores transactions for money laundering risk, replacing a rules-based system that raised many false alarms.

HSBC says it is finding two to four times more financial crime than before, with 60% fewer false positive cases (customers no longer called unnecessarily about legitimate activity).

SOURCE [71]

YOUR NOTES

Which of these would fit your business, and who would own it?

PART 3

Companies your size

$10M to $100M operators who put AI to work with small teams: what they built, who checks it, and what changed.

Operators doing it with small teams

FREIGHT BROKERAGE AND 3PL (INDIANAPOLIS), ~$60M REVENUE

InTek Logistics

The CEO did some of the accounting transactions himself to see where time went and found one customer's accessorial-charge invoices needed several documents (charge sheet, approval backup, invoice) combined by hand. An AI agent from Ventus AI now builds each invoice, gathers and combines the documents into one file and emails it, working through screens rather than system integrations.

CEO Rick LaGore says each packet took "about 7 minutes, 8 minutes" by hand; the agent now does "like 150 of these transactions in... like 3 minutes". Staff time moved to customer-facing work.

> Do the tedious job yourself for a day; the steps that make you wince are where an agent pays back fastest.

SOURCE [72]

FAST-FASHION MENSWEAR BRAND (INDIA), ~$59M REVENUE

Snitch

Snitch built AI agents that file and settle claims with its courier partners for lost or damaged shipments. The agent sends the first response on about 80% of cases and settles most without a person; Snitch builds 90-95% of its AI in-house.

CTO Maruthy Ramgandhi says 85-90% of claims are resolved with no human involvement, claim handling fell from about 35 minutes to under one minute, and the agents recover approximately 45 lakh rupees (about $50,000) a month.

> Money you are owed but rarely chase, like courier claims, is a clean first agent because the rules are clear and the payoff is cash.

SOURCE [73]

BANKRUPTCY CASE ADMINISTRATION FIRM, ~$64M REVENUE, ~330-400 STAFF

Stretto

After a Chapter 11 filing, creditors flood a call centre with questions. Stretto built 'Conductor', which ingests court filings within minutes and answers creditor questions 24/7 with a cited source for every answer, using its own ranking of which documents are most authoritative and recent.

A Stretto leader says Conductor "helped us reduce our call center cost by 90%, the operation cost by 90%", replacing calls that averaged about 9 minutes with instant answers.

> When the same questions arrive in waves, a cited answer engine over your own documents can replace most of the call centre.

SOURCE [74]

INSURANCE AGENCY (WISCONSIN), ~$38M-$75M REVENUE, ~300-500 STAFF

M3 Insurance

M3's employee-benefits team uses an AI tool (Stitch Studio) to build the health plan comparison documents clients need before renewal, treating the AI like a junior colleague whose work is always checked. Carriers sometimes return proposals only about three days before the client meeting.

M3 reports a comparison that took 90-120 minutes to build now takes about 15 minutes; the freed time goes to client explanation and to seniors coaching junior staff.

> Treat AI as a junior colleague whose drafts are always reviewed, and expect a longer ramp than the first few weeks suggest.

SOURCE [7]

YOUR NOTES

Pick the case on this page closest to your business.

CLOSEST TO MY BUSINESS: _____

THE JOB IT WOULD TAKE OFF US: _____

WHO WOULD CHECK THE OUTPUT: _____

HOW WE WOULD KNOW IT WORKED: _____

More operators your size

DEFENSE PARTS RESELLER, $20M REVENUE, 5 STAFF

Echelon Supply & Service

Echelon bids on US Defense Logistics Agency solicitations. A platform from SalesPatriot now reads new solicitations, finds the right manufacturers and suppliers, sends tailored quote-request emails, and parses supplier replies for the sales team, who still review quotes and submit the bids.

SalesPatriot reports bid volume up 283% over a three-month deployment, with 3.3x more supplier emails a day and 2.3x more vendor quotes a week, with the team unchanged at five people.

> For a small team, AI's biggest win can be volume: answering many more opportunities with the same people.

SOURCE [75]

AIRCRAFT MODIFICATION AND MAINTENANCE FIRM, ~250 STAFF, ~$56-75M REVENUE

Field Aerospace

Field's business development staff upload government solicitations to an internal web app. Workflows on n8n split them into sections, draft responses from an approved library including past performance, build a requirements matrix of 'shall/must/will' statements, and send labelled AI drafts to Microsoft Teams for people to edit.

n8n reports a proposal draft that took about two weeks of work from three to four people now reaches an 80% starting point in roughly 25 minutes; the requirements matrix takes 15-20 minutes instead of hours.

> Proposals built from your own past answers are a fast win: let AI produce the 80% draft and keep experts on pricing and strategy.

SOURCE [76]

OUTSOURCED ACCOUNTING FIRM FOR STARTUPS, ~$33M REVENUE, ~120-180 STAFF

Countsy

Countsy's AP team codes up to 3,500 client invoices a month. It adopted Vic.ai, which predicts the accounting codes (account, department, location) for each invoice so accountants review rather than type; new clients are onboarded to the tool in about two weeks.

Vic.ai reports, quoting Countsy, an 84% decrease in processing time per invoice (from 4-5 minutes to under 2), 95% coding accuracy and 78% of invoices running with no human input, freeing hundreds of hours a month from a ~300-hour manual baseline.

> Invoice coding is a proven, low-risk place to start in finance; measure minutes per invoice before and after.

SOURCE [77]

VENDOR-RISK SOFTWARE AND SERVICES FIRM, ~$36-42M REVENUE

Venminder (Ncontracts)

Venminder's analysts extract key terms (dates, auto-renewal, clauses) from client contracts. With Caylent it added Claude on AWS to search documents in plain language, read scans and pull the fields into its required formats, processing documents in parallel.

Caylent reports extraction accuracy rose from 60% to 86%+ (against a 75% target), about 70% of analysts' time was freed, and the contract review backlog fell from 65 days to 3 days within 4 days of adding generative AI.

> Set an accuracy target before you start; an 86% first pass with human checks can clear a backlog that hiring never would.

SOURCE [8]

YOUR NOTES

Pick the case on this page closest to your business.

CLOSEST TO MY BUSINESS: _____

THE JOB IT WOULD TAKE OFF US: _____

WHO WOULD CHECK THE OUTPUT: _____

HOW WE WOULD KNOW IT WORKED: _____

More operators your size

AMERICAN-MADE APPAREL AND GEAR BRAND SELLING DIRECT, ~$10-25M REVENUE, ~50 STAFF

Origin USA

Customer service suspected one light-wash denim fit inconsistently. Origin's marketing lead gave an AI tool the returns, reviews and production tech packs together; it traced the problem to the wash process affecting shrinkage and recommended tighter tolerances on specific measurements, checks against spec and new site messaging. The product team verified the findings before changing anything.

Projection: Origin's Justin Parker says the AI's readout, based on return rates, exchanges and return shipping for that wash, estimated the fix would save "like 150 grand a year". Not a measured saving.

> Put returns, reviews and product specs in one place and ask AI why something is going wrong; then have the experts verify before acting.

SOURCE [78]

UK SUPPLEMENTS BRAND SELLING DIRECT (ORKLA-OWNED), ~GBP 50M TURNOVER, 50-SEAT CONTACT CENTRE

Healthspan

Healthspan's 'Product Professor' chat assistant answers customer questions about which vitamins and supplements suit them from the company's product knowledge base, and always offers a handoff to the contact centre at the end of a chat. Questions it cannot answer are passed to marketing as signals of products customers want.

Head of IT Rob King reports 90% of product-related customer queries in chat are now fully automated.

> A well-fed product assistant can take most pre-sale questions, and the questions it can't answer tell you what to stock next.

SOURCE [3]

NONPROFIT FOOD BANK (TEXAS), ~$38M REVENUE INCL. DONATED FOOD (FY2024 FORM 990), 46 STAFF

Montgomery County Food Bank

A small team used Claude, including Claude in Excel, to analyse driver and warehouse shifts against delivery requirements, compare a partner agency's purchases with the food bank's at-cost prices, and benchmark against peer food banks using public filings instead of hiring consultants.

Anthropic reports biweekly driver and warehouse overtime fell from 99 hours to 11 after the shift analysis (CEO Scott Burns credits it "in large part" to the Claude-supported analysis), and one large partner agency's food costs fell 40%.

> You do not need an agent to get value; a manager using AI to analyse shifts or purchasing data can cut costs within weeks.

SOURCE [79]

WAXING SALON FRANCHISE (ARIZONA), ~$13.5M REVENUE, ~94 HQ STAFF

Hello Sugar

Hello Sugar wanted to double its number of salons without growing its reception team. It combined scripted flows with generative AI answers in Zendesk so customer messages about bookings, pricing and services get answered automatically, with humans taking the rest.

Zendesk reports a 66% automation rate, $14K in monthly savings and 200 agent hours saved a month.

> Mix fixed scripts for the common paths with AI for the open questions, and you can grow locations without growing the front desk.

SOURCE [80]

YOUR NOTES

Pick the case on this page closest to your business.

CLOSEST TO MY BUSINESS: _____

THE JOB IT WOULD TAKE OFF US: _____

WHO WOULD CHECK THE OUTPUT: _____

HOW WE WOULD KNOW IT WORKED: _____

PART 4

Your plan

Worksheets to turn what you have read into a first project: list the ideas, find where the hours go, put a number on each, test it on your own work and roll it out in stages.

List the ideas on the table

Most companies already have a dozen AI ideas floating around. Write down what people have suggested, what it would need to do to be useful, and tick the ones worth a closer look. These become your shortlist.

AREA	THE IDEA YOU'VE HEARD	TO BE USEFUL IT SHOULD	SHORTLIST?
Customer support	"Put a chatbot on the website."	Answer from your own help content, and hand the hard cases to a person with the conversation attached.	☐
Finance and AR	"Can AI do our reconciliations?"	Match payments and remittances to invoices, flag short-pays and deductions, post only after review.	☐
Orders and fulfillment	"Where is my order?" eats the day.	Pull status from every system and draft the reply, or answer it directly.	☐
Sales	"Have AI write our follow-ups."	Draft follow-ups from call notes and CRM history, in your voice, for a rep to send.	☐
Reporting	"I want to ask the data questions."	Answer plain-English questions over live company data, read-only, with the numbers checked.	☐
Company knowledge	"Nobody can find anything."	Search policies, documents and past answers, and show the source behind every answer.	☐
Operations	"People copy data between systems."	Move records between the ERP, CRM and spreadsheets, with a log and a person approving changes.	☐
Marketing	"Use AI to make more content."	Draft on-brand content and pull performance from your ad and analytics platforms.	☐
Your own idea			☐

IN ONE SENTENCE, WHICH IDEA WOULD MAKE THE BIGGEST DIFFERENCE IF IT WORKED?

Map where the hours actually go

Estimate a normal week across the team, and count everyone who touches the work. Most teams find the total is far higher than they guessed, because the work arrives in small pieces all day.

JOB	WHO DOES IT NOW	HOURS / WEEK	AI COULD DRAFT IT?	NEEDS SIGN-OFF?
Customer emails and tickets			☐	☐
Order-status questions			☐	☐
Matching payments to invoices			☐	☐
Copying data between systems			☐	☐
Building reports			☐	☐
Chasing people for updates			☐	☐
Follow-ups, quotes and proposals			☐	☐
Finding documents and answers			☐	☐
Your own row			☐	☐

Total hours a week on repetitive work

Multiply by 50 for the year. That's the time a first AI project can start giving back.

TIP

Tick "Needs sign-off?" for anything that spends money, changes a record or speaks for the company. Those jobs can still move to AI, as drafts a person approves.

WATCH FOR

Work nobody owns. If "Who does it now" is blank, that's usually where errors and missed money hide.

AI as seats, and AI as a plan

The difference isn't the model. It's whether the AI is connected to your systems, aimed at one process where a number moves, and built with the people who will use it.

AI AS SEATS

Everyone gets a chat window

Licenses go out, an announcement goes round, and usage peaks in the first week.

Nothing is connected

It can't see the ERP, CRM or helpdesk, so people paste data in by hand.

Easy wins first

The first projects are chatbots and summaries because they're easy to build, not because they move a number.

Edge cases break trust

It handles most cases and fumbles the rest, so people stop relying on it.

No way to tell if it paid off

Nobody measured before, so nobody can say what changed.

AI AS A PLAN

One process at a time

A single, frequent piece of work where a number actually moves.

Connected to your systems

It reads the data it needs from the source, read-only first.

Built with the people who use it

The team doing the work shapes it, so it fits how they really work.

The last 10 percent goes to a person

Anything it can't handle is handed over with the context attached.

Measured before and after

You know the hours, errors or response times it changed.

WHAT IT SOUNDS LIKE, IN PRACTICE

ILLUSTRATION

How many orders shipped late last week, and why?

38 orders shipped late. 22 were waiting on one SKU that was out of stock at the warehouse, 11 were held for address checks, and 5 were wholesale orders entered after the cutoff.

Draft a note to the warehouse about the SKU and list the wholesale ones.

Done. The note is waiting for your OK, and the five wholesale orders are listed with their customers and entry times.

Book a free discovery session →

Scan or click to pick a time with a senior AI architect.

Put a number on it

Time is only half the case. The other half is what slips through when nobody has time to look: errors, slow responses and deductions nobody disputed. Use last quarter as your guide.

WHERE VALUE LEAKS	HOW TO ESTIMATE IT	LAST QUARTER
Rework and errors	Orders, invoices or records fixed after the fact, times the time and cost to fix each one.	
Slow responses	Tickets, quotes or leads answered after a day or more. Estimate the share that went cold.	
Deductions and chargebacks	Short-pays, retailer deductions and chargebacks nobody disputed in time.	
Data entry between systems	Hours spent copying records, multiplied by a loaded hourly cost.	
Decisions waiting on reports	Days spent assembling numbers before a decision could be made.	

Estimated value per quarter

Add it to the hours from 1.2. Together, that's the case for your first project.

VALUE BEYOND THE NUMBERS

01

Growth without hiring at the same pace

Volume grows, and the repetitive work doesn't grow with it.

02

Faster answers

Customers, retailers and your own team get answers the same day.

03

Knowledge that stays

How the work gets done lives in a system, not in one person's head.

Check your plan against the six pitfalls

Ask these questions of any AI plan, whoever builds it. If the answer to the first one is "we'll start with a chatbot because it's quick", stop and think again.

AREA	QUESTION TO ASK	WHAT GOOD LOOKS LIKE
Use case	Is the first project the most useful one, or just the easiest?	One piece of one process a team does every day, where a number moves.
Adoption	Who will use it, and were they involved?	The people doing the work help define it and test it before launch.
Edge cases	What happens to the cases it can't handle?	They go to a person with the context attached. It never guesses.
Systems	Is it connected to the data it needs?	Read-only access to the source systems first, not a copy-and-paste workflow.
IT and governance	Where does it run, what can it touch, and who approves?	Agreed with IT before building: access by role, data kept in your environment.
Measurement	How will you know it paid off?	A number measured before and after, with a first investment small enough to prove it.

The rule of thumb: the first project should be small enough to prove and useful enough to matter. If it's only one of those, pick again.



Book a free discovery session →

Scan or click to pick a time with a senior AI architect.

Build a realistic test from your own work

A demo on clean sample data proves nothing. Before anything goes live, collect 15 to 20 real situations from the last few months, including the messy ones, and run them in read-only or draft mode first.

TYPE OF SITUATION	EXAMPLE FROM A TYPICAL COMPANY	YOUR EXAMPLE
The everyday case	A customer asks where their order is, and the order is on time.	
Conflicting data	The ERP and the warehouse show different stock for the same SKU.	
Missing information	An invoice arrives with no PO number on it.	
An exception	A wholesale customer asks for terms outside your usual policy.	
Something sensitive	A complaint mentions a health issue or a legal threat.	
Out of bounds	Someone asks it to change a price or issue a large refund.	
A vague request	"Can you sort out the thing with that order from last week?"	

Keep this list. It becomes the test set every new version has to pass before it goes live, so a change that breaks something shows up in a test run, not in front of a customer.

Score the results

Run your test situations and score each area from 1 to 5. Anything below a 4 on accuracy or safety means it isn't ready to act on its own yet, however impressive it looks.

AREA	WHAT TO CHECK	SCORE 1-5
Accuracy	Are the numbers right? Amounts, dates, order and invoice details match the source systems.	1 · 2 · 3 · 4 · 5
Completeness	Did it find everything? Every open item, every exception.	1 · 2 · 3 · 4 · 5
Judgement	Did it prioritise sensibly and ask when the request was unclear?	1 · 2 · 3 · 4 · 5
Draft quality	Would your team send its drafts with little or no editing?	1 · 2 · 3 · 4 · 5
Safety	Did it stop for approval every time it should, and stay inside its limits?	1 · 2 · 3 · 4 · 5
Traceability	Can you see what it did, which data it used and why?	1 · 2 · 3 · 4 · 5
Speed	Is it faster than doing it by hand, including the time to check its work?	1 · 2 · 3 · 4 · 5

Total score (out of 35)

28 or more with no area below 4: ready for drafts with approval. Below that: keep it read-only and fix the gaps.

Know what a first version includes

Most AI vendors stay vague about the first release, and you find out what's missing after you've paid. You should know the scope before you start. This is what a solid first version covers.

01

One use case, done properly

The single process you picked in step 1, end to end.

"Handle order-status questions."

02

Connected to your systems

Read-only access to the ERP, CRM or helpdesk data it needs.

"What's the status of order 48211?"

03

Inside the tools you already use

Your helpdesk, CRM, Slack or Teams, not a new screen to learn.

"Draft the reply in the ticket."

04

Approvals on anything that writes

Changes, refunds and customer messages wait for a person.

"What needs my approval?"

05

A test set from your own cases

The situations from 2.2, run before and after every change.

"Run the test set."

06

Logs, training and handover

Every action written down, your team trained, code and docs handed over.

"Show me what it did yesterday."

NOT IN A FIRST VERSION, ON PURPOSE

Company-wide rollouts, fully automatic actions and anything that replaces a core system. Those come once the first version has earned trust on your real data.

KEEP WHAT ALREADY WORKS

If a tool already does a job well, AI connects to it rather than rebuilding it. It takes on the work between systems, not the systems themselves.

Roll it out in four stages

Each stage is live and useful on its own, and you decide when the next one starts. If something isn't right, you've lost a week of one stage, not the whole project.

STAGE 0

Map it

The use case, data sources, approval rules and security, agreed with IT. A design you can hold the builder to, before any code.

DONE WHEN

You've signed off the design, the access list and the approval rules.

STAGE 1

Prove it

A working version on your own data, read-only. It can see what it needs and change nothing.

DONE WHEN

Its answers match your systems on every test situation from 2.2.

STAGE 2

Drafts with approval

It drafts the replies, updates and records, and a person approves each one.

DONE WHEN

Your team approves most drafts without edits, scoring 28+ on 2.3.

STAGE 3

Trusted actions

Low-risk actions inside agreed limits run on their own. Everything else still waits for a person.

DONE WHEN

A month of clean logs, and you've chosen which actions to trust.

Trust grows with evidence, never with a switch.



FIG. 1 · TYPICAL TIMING FOR ONE USE CASE. MORE SYSTEMS TAKE LONGER AT STAGE 0.

Book a free discovery session →

Scan or click to pick a time with a senior AI architect.

Write your approval rules

Decide in advance what AI may do on its own, what needs a person, and what it should never touch. We've filled in a sensible starting point. Change any row to suit how you work.

ACTION	ON ITS OWN	NEEDS APPROVAL	NEVER	YOUR LIMIT OR NOTE
Answer questions, build reports	■	☐	☐	
Create tasks and reminders	■	☐	☐	
Draft a customer reply	☐	■	☐	
Update a CRM or ERP record	☐	■	☐	
Send an email for the company	☐	■	☐	
Issue a refund or credit	☐	■	☐	
Change prices or terms	☐	☐	■	
Pay invoices or move money	☐	☐	■	
Delete records	☐	☐	■	
Your own row	☐	☐	☐	

Start strict, loosen with evidence. Move an action from "needs approval" to "on its own" only after the log shows it was approved unchanged many times in a row.

Measure whether it's working

Set checkpoints now, using your answers from step 1, so success is something you measure rather than something you feel. Write your target next to each one.

WHEN	WHAT TO MEASURE	YOUR TARGET
Week 1	Its answers match what your team would have found, on the test set and on live questions.	
Week 4	Share of drafts approved without edits. Hours a week back from the list in 1.2.	
Week 8	Errors, response times or deductions versus last quarter. Who uses it every day.	
Week 12	Value from 1.4 versus your estimate. The next use case, picked with real numbers.	

SIGNS IT'S READY FOR THE NEXT STAGE

- ✓ Your team checks its work out of habit, not because they've found mistakes.
- ✓ The action log has no surprises in it.
- ✓ People have started asking it things they used to ask a colleague.

Sources

Every result in this report is as the company, its executive or its vendor reported it in public. Check the source before you quote a number.

1. **Johnson & Johnson** · Jim Swanson (J&J CIO), Reuters Momentum AI conference, April 2026, via Yahoo Finance
   https://finance.yahoo.com/sectors/healthcare/articles/j-j-sees-ai-halving-150527193.html
2. **Thermo Fisher Scientific** · Harry Look (Thermo Fisher), Dataiku event, ~April 2025, 00:11:17-00:13:10
   https://www.youtube.com/watch?v=T1bKUIJpHN4&t=677s
3. **Healthspan** · Rob King (Healthspan Head of IT), Talkative webinar, March 2026, ~00:21:17-00:22:06
   https://www.youtube.com/watch?v=iLOrwKAK5Mg&t=1277s
4. **Healthspan** · Rob King (Healthspan Head of IT), Talkative webinar, March 2026, 00:15:30-00:16:40
   https://www.youtube.com/watch?v=iLOrwKAK5Mg&t=930s
5. **Salesforce** · Bernard Slowey and Joe Inzerillo (Salesforce), VentureBeat, July 2025
   https://venturebeat.com/ai/salesforce-used-ai-to-cut-support-load-by-5-but-the-real-win-was-teaching-bots-to-say-im-sorry/
6. **Amazon (AWS)** · Financial Times report via Engadget, Feb 2026
   https://engadget.com/ai/13-hour-aws-outage-reportedly-caused-by-amazons-own-ai-tools-170930190.html
7. **M3 Insurance** · Amber Seitz (M3 Insurance), Stitch Studio podcast, Oct 2026, ~00:16:18
   https://www.youtube.com/watch?v=qJ2JT1J7RDM&t=978s
8. **Venminder (Ncontracts)** · Caylent case study
   https://caylent.com/case-study/venminder-generative-ai
9. **JPMorgan Chase** · Derek Waldron (JPMorgan Chief Analytics Officer), VentureBeat Beyond the Pilot video, 2025, 00:05:40-00:06:56 and 00:32:50-00:33:30
   https://www.youtube.com/watch?v=rLxJzeRGzV8&t=340s
10. **BBVA** · Antonio Bravo (BBVA Global Head of AI Transformation), OpenAI customer talk, 2025, 00:12:40-00:14:05; OpenAI BBVA customer story
    https://www.youtube.com/watch?v=UNJSk90Lz1c&t=760s
11. **Bank of America** · Bank of America press release, April 2025
    https://newsroom.bankofamerica.com/content/newsroom/press-releases/2025/04/ai-adoption-by-bofa-s-global-workforce-improves-productivity--cl.html
12. **Indeed (Recruit Holdings)** · Michael Redding and Jeff Davis (Indeed), DX Annual session, DX newsletter, 2026
    https://newsletter.getdx.com/p/2x-the-power-users-how-structured
13. **IBM** · Nickle LaMoreaux (IBM CHRO), Fortune, July 2024; IBM AskHR case study, 2025
    https://fortune.com/2024/07/12/ibm-chro-ai-chatbot-rollout-strategy-lessons/
14. **Samsung Electronics** · Bloomberg via The Star, May 2023 (ban); OpenAI announcement, June 2026 (rollout); Seoul Economic Daily video, May 2026 (training gate)
    https://openai.com/index/samsung-electronics-chatgpt-codex-deployment
15. **Goldman Sachs** · Bloomberg report, Nov 2026, via Dealroom News
    https://dealroom.co/news/169162-goldman-restricts-claude-code-github-for-traders-and-bankers/
16. **Palo Alto Networks** · Meerah Rajavel (Palo Alto Networks CIO), Fortune, June 2025
    https://fortune.com/2025/06/18/palo-alto-networks-cio-gen-ai/
17. **Airbus** · Guillaume Faury (Airbus CEO), Adopt AI Grand Palais interview, 2025, 00:01:20-00:02:05
    https://www.youtube.com/watch?v=UkCs3Hd7OUU&t=80s
18. **Santander** · Inaki Bernal (Santander Chief AI Officer), Fintech NerdCon with OpenAI, ~late 2025, 00:04:33-00:04:53 and 00:14:36-00:15:15
    https://www.youtube.com/watch?v=3x5aa0ycnCc&t=873s
19. **Philip Morris International** · PMI AI lead, Deloitte 'Generative Age' interview, 2026, 00:05:41-00:06:48 and 00:11:00-00:11:24
    https://www.youtube.com/watch?v=H21-v5BEzGM&t=341s
20. **Accenture** · 404 Media report on a leaked Accenture memo, Oct 2026, via news recaps (The SI Brief)
    https://www.youtube.com/watch?v=1nCuZ-NlCEg&t=22s
21. **AT&T** · Andy Markus (AT&T Chief Data Officer), CNBC AI Forum video, 2026, 00:03:20-00:04:16; VentureBeat, 2026
    https://www.youtube.com/watch?v=Q7496Cix9PQ&t=196s
22. **Cisco** · Jeetu Patel (Cisco President), CXOTalk, Oct 2026, 00:10:55-00:12:31
    https://www.youtube.com/watch?v=ZCN5hkJJH9c&t=655s
23. **Air Liquide** · Amaury de Menonville (Air Liquide), BoursoLive shareholder webinar, June 2026, published on airliquide.com
    https://www.airliquide.com/stories/shareholding/discover-boursolive-webinar-involving-individual-shareholders-promises-ai
24. **Citigroup** · David Griffiths (Citi CTO), American Banker, July 2025; Jane Fraser, Q3 2025 earnings call via Reuters
    https://www.americanbanker.com/news/citi-is-rolling-out-agentic-ai-to-its-40-000-developers
25. **DBS Group** · DBS, via Business Times reporting summarised by FCC Singapore and Fintech News Singapore, Feb 2026; DBS on Frontier Enterprise
    https://www.frontier-enterprise.com/dbss-approach-to-ai-at-scale/
26. **Meta Platforms** · Susan Li (Meta CFO), Q4 2025 earnings call via Computerworld; Meta developer productivity team, Develocity talk, 00:20:41-00:21:57
    https://www.youtube.com/watch?v=1OzxYK2-qsI&t=1230s
27. **ABB** · Alec Joannou (ABB Global CIO), CONFARE interview, 00:16:48-00:17:16; ABB 'AI in the workplace' page, 2026
    https://www.youtube.com/watch?v=G7yb2JEGUJw&t=1008s
28. **IBM** · Jim Kavanaugh (IBM CFO), Q2 2025 earnings call, via Next Platform, July 2025; IBM Client Zero case study, 2026
    https://www.nextplatform.com/2025/07/24/for-now-ai-helps-ibms-bottom-line-more-than-its-top-line/
29. **ExxonMobil** · ExxonMobil executive with SAP, SAP Sapphire Orlando 2026, 00:05:54-00:06:21
    https://www.youtube.com/watch?v=7VdD-7Qpv_4&t=354s
30. **Virgin Atlantic** · Tom Barber (Virgin Atlantic Head of BI), Databricks conference talk, 00:06:48-00:10:44
    https://www.youtube.com/watch?v=coJor_d5kGw&t=408s
31. **Eaton** · Eaton director of analytics, Dataiku Everyday AI conference, 00:03:18-00:10:57
    https://www.youtube.com/watch?v=KqIt3LK4CYU&t=598s
32. **Prosus** · Prosus, 'The Coming Age of AI Colleagues' report, 2026
    https://www.prosus.com/~/media/Files/P/prosus-corp-v2/documents/the-coming-age-of-ai-colleagues.pdf
33. **SoftBank Corp.** · SoftBank News, Dec 2025 (2.5M agents); Business+IT (SB Creative), Oct 2026 (6,000+ kept)
    https://www.softbank.jp/en/sbnews/entry/20251217_01
34. **Oracle** · Jae Evans and Clay Magouyrk (Oracle), internal town hall reported by The Next Web, Sept 2026
    https://thenextweb.com/news/oracle-internal-ai-rollout-town-hall

Sources (continued)

Every result in this report is as the company, its executive or its vendor reported it in public. Check the source before you quote a number.

35. **Machinecraft** · Rushabh Doshi (Machinecraft), AI Engineer talk, ~2026, 00:05:50-00:06:41 and 00:08:21-00:09:00
https://www.youtube.com/watch?v=jtzh-GBXBWc&t=350s
36. **Edward Don & Company** · Tim Walter (Edward Don CIO), Technology Executives Club case study, Feb 2026, ~00:16:58-00:17:40
https://www.youtube.com/watch?v=3qky7K1fFCQ&t=1018s
37. **BioIVT** · LaRae Bennett (BioIVT), Salesforce customer story Q&A, May 2026
https://www.salesforce.com/customer-stories/bioivt-larae-bennett-interview/
38. **Uber for Business** · Salesforce customer story, Aug 2026
https://www.salesforce.com/customer-stories/uber/agentic-outreach/
39. **SMBC (Sumitomo Mitsui Banking Corporation)** · SMFG press release (Japanese), April 2026; Nikkei via iThome, 2026
https://www.smfg.co.jp/news/pdf/j20260430_01.pdf
40. **Ridge** · Sean Frank (Ridge CEO), Operators Podcast, March 2026, 00:23:17-00:23:45; cost per ticket from Operators Podcast, April 2026, 00:25:38
https://www.youtube.com/watch?v=iOnG0kyxIII&t=1397s
41. **T-Mobile US** · Julianne Roberson (T-Mobile), AI Agent Conference, reported by Entrepreneur, May 2026
https://www.entrepreneur.com/business-news/t-mobile-director-says-half-t-mobile-customer-calls-ai
42. **Verizon** · CX Today, April 2025; Shankar Arumugavelu (Verizon), Metis Strategy Technovation podcast, 00:09:56-00:10:38
https://www.cxtoday.com/contact-center/verizon-teams-up-with-google-to-give-28000-customer-support-reps-a-personal-research-assistant/
43. **Mercari US** · Zapier customer story, 2026
https://zapier.com/customer-stories/mercari
44. **Kogan.com** · Kogan.com speaker, Salesforce video, 2026, 00:05:15-00:07:41
https://www.youtube.com/watch?v=ZDDjKo-POA8&t=315s
45. **Microsoft** · Cory Hrncirik (Microsoft modern finance leader), CFO Executive Briefing with Alpha Variance Solutions, 00:30:11-00:31:03
https://www.youtube.com/watch?v=yFj4AEj3x6E&t=1811s
46. **Pamarco Technologies** · Glyn Wilson (Pamarco CIO), Microsoft customer story, 2026
https://www.microsoft.com/en/customers/story/27289-pamarco-technologies-dynamics-365-finance
47. **Diesel Direct** · Andrew Turkish (Diesel Direct VP and Corporate Controller), Vic.ai case study, Feb 2025
https://www.vic.ai/resources/case-studies/diesel-direct-fuels-growth-with-ai-powered-accounting
48. **Hewlett Packard Enterprise** · Deloitte case article, 2025; Fortune, Feb 2026
https://www.deloitte.com/us/en/services/consulting/articles/hpe-agentic-ai-collaboration.html
49. **DoorDash** · Gordon Lee (DoorDash CAO), Center for Digital Transformation (UC Irvine) event, Nov 2026, 00:11:41-00:12:36
https://www.youtube.com/watch?v=7N8OLyEd440&t=701s
50. **Millcraft Paper Company** · Travis Mlakar (Millcraft President), Wholesale Change Show, Distribution Strategy Group, May 2024, 00:14:30-00:15:40
https://www.youtube.com/watch?v=LBsIGKOsgks&t=870s
51. **C.H. Robinson** · LangChain customer story, March 2025
https://www.langchain.com/blog/customers-chrobinson
52. **Choco (for food distributors)** · OpenAI customer story, April 2026
https://openai.com/index/choco/
53. **Dude Wipes** · Zack Means (Dude Wipes Senior Operations Analyst), Make customer story, Dec 2025
https://www.make.com/success-stories/how-make-helped-dude-wipes-build-a-culture-that-fuels-rapid-growth
54. **Glamnetic** · Mia Chapa (Glamnetic VP CX), Gorgias customer story, Dec 2024
https://www.gorgias.com/customers/glamnetic-support
55. **IBM** · IBM AskIT transformation case study, 2026
https://www.ibm.com/case-studies/ibm-transformation/ask-it
56. **Broadcom** · Moveworks customer story with Broadcom CIO Andy Nallappan, 2025
https://www.moveworks.com/us/en/customers/broadcom-integrates-it-knowledge-base-with-moveworks-ai
57. **Commonwealth Bank of Australia** · Microsoft Source Asia, July 2025
https://news.microsoft.com/source/asia/2025/07/01/commonwealth-banks-chatit-helps-employees-fix-it-issues-in-seconds-using-generative-ai/
58. **Schneider Electric** · Anna Gawlicka (Schneider Electric), company podcast EP28, 00:06:32-00:07:33
https://www.youtube.com/watch?v=guf98AM3ofg&t=392s
59. **Uber** · Uber Engineering blog, Aug 2025
https://www.uber.com/us/en/blog/ureview/
60. **Stripe** · Stripe engineering blog, 'Minions' Part 2, Feb 2026
https://stripe.dev/blog/minions-stripes-one-shot-end-to-end-coding-agents-part-2
61. **Airbus** · Catherine Jestin (Airbus CIO), Paris Air Show, reported by IMechE Professional Engineering, 2025
https://www.imeche.org/news/news-article/ai-won-t-take-your-job-but-someone-who-knows-ai-will-airbus-boss-warns-paris-air-show-audience
62. **Spotify** · Spotify Engineering blog, Honk Part 1, Nov 2025
https://engineering.atspotify.com/2025/11/spotifys-background-coding-agent-part-1
63. **Gumroad** · Sahil Lavingia (Gumroad CEO), How I AI podcast, ~Nov 2025, 00:02:40-00:02:55
https://www.youtube.com/watch?v=KVZ3vMx_aJ4&t=160s
64. **Procter & Gamble** · Seth Cohen (P&G CIO), Metis Strategy Technovation podcast, 2025, 00:24:20-00:26:20
https://www.youtube.com/watch?v=16y0JDXbnw4&t=1460s
65. **Toyota** · Kohdai Gotoh (Toyota), Google Cloud blog, Dec 2024
https://cloud.google.com/blog/topics/hybrid-cloud/toyota-ai-platform-manufacturing-efficiency
66. **GE Aerospace** · GE Aerospace speaker, Palantir AIPCon 9, March 2026, 00:05:16-00:05:40
https://www.youtube.com/watch?v=ipdtAHQj0l0&t=316s
67. **Unilever** · Unilever newsroom, 2026
https://www.unilever.com/news/news-search/2026/building-resilience-driving-growth-inside-unilevers-aipowered-supply-chain/
68. **Cascade Engineering** · Christina Keller (Cascade Engineering CEO), MIT Sloan Management Review 'Me, Myself, and AI', Dec 2025, ~00:05:40-00:06:20
https://www.youtube.com/watch?v=_AgepJx_W1M&t=340s

Sources (continued)

Every result in this report is as the company, its executive or its vendor reported it in public. Check the source before you quote a number.

69. **Morgan Stanley** · Jeff McMillan (Morgan Stanley), TechXchange 2025 New York talk, 00:04:48-00:05:08
https://www.youtube.com/watch?v=qFXwi8lZr9Q&t=288s

70. **Nestle** · Nestle spokesperson, FoodNavigator, June 2025; Chris Wright (Nestle CIO), Fortune, Jan 2026
https://www.foodnavigator.com/Article/2025/06/17/what-are-nestle-unilever-kraft-heinz-and-mondelez-doing-with-ai

71. **HSBC** · HSBC 'Harnessing the power of AI to fight financial crime', company article, 2023-2024
https://www.hsbc.com/news-and-views/views/hsbc-views/harnessing-the-power-of-ai-to-fight-financial-crime

72. **InTek Logistics** · Rick LaGore (InTek Logistics CEO), InTek x Ventus AI podcast, Aug 2025, 00:11:29-00:12:40
https://www.youtube.com/watch?v=_eF34CVIaxY&t=689s

73. **Snitch** · Maruthy Ramgandhi (Snitch CTO), Pulse AI x D2C Insider talk, Aug 2026, 00:18:30-00:19:30
https://www.youtube.com/watch?v=FA-z2rK1or0&t=1110s

74. **Stretto** · Stretto, AWS Events talk, 00:04:20-00:04:45
https://www.youtube.com/watch?v=EF7iowMYTw8&t=260s

75. **Echelon Supply & Service** · SalesPatriot case study, 2026
https://www.salespatriot.com/case-studies/echelon-supply-service-salespatriot

76. **Field Aerospace** · n8n customer case study, Oct 2026
https://n8n.io/case-studies/field-aerospace/

77. **Countsy** · Countsy Director of Technology, Vic.ai case study, Jan 2025
https://www.vic.ai/resources/case-studies/countsy-case-study

78. **Origin USA** · Justin Parker (Origin), Klaviyo Originals, Oct 2026, ~00:09:00-00:10:15; The DTC Podcast, June 2026, 00:17:02-00:18:39
https://www.youtube.com/watch?v=rcadl9TmGqk&t=540s

79. **Montgomery County Food Bank** · Scott Burns (Montgomery County Food Bank CEO), Anthropic customer story
https://claude.com/customers/montgomery-county-food-bank

80. **Hello Sugar** · Austin Towns (Hello Sugar CTO), Zendesk customer story, Aug 2024
https://www.zendesk.com/customer/hello-sugar/

Pick one project.
Measure it. Then the next.

Every company in this report started small: one workflow, one team, one number to move. Use the worksheets to pick yours, write down today's number, and give it a few weeks.

01

Pick the work

Use the hours map and the value worksheet to choose one project with real hours or money behind it.

02

Keep a person on it

Let AI draft and a person approve, with limits on what it can touch, until the numbers earn more trust.

03

Measure honestly

Compare before and after, count the running cost, and keep only what pays back.

AiSolv builds AI into the operations of $10M to $100M companies: ERP, CRM, helpdesk and data. If you want help with your first project, you'll find us at **aisolv.io**.