900 LABS

The AI Execution Gap

Why about 95% of Enterprise AI Fails (or Stalls), and What the Top 5% Do Differently to Reach Measurable ROI

900 LABS · MARCH 2026

Executive Summary

> Enterprise AI is not stalling because the models are weak. It stalls because the execution stack is weak: data, operating model, and measurement.

"Failure" depends on the denominator. MIT Project NANDA reports 95% of organizations see zero measurable P&L return from enterprise GenAI, while 5% of integrated pilots extract meaningful value, despite an estimated $30 billion to $40 billion invested; the report is interview-derived and labeled directional.[1] S&P Global and 451 Research find organizations scrap 46% of AI proof-of-concepts before production, and the share abandoning the majority of AI initiatives rose from 17% to 42% year over year.[2] Gartner reports more than half of GenAI projects were abandoned after proof-of-concept by the end of 2025.[3]

95%	$632B	92%
ZERO MEASURABLE P&L RETURN	AI SPEND PROJECTED BY 2028	CITE PEOPLE, NOT TECH, AS BARRIER

The gap is expensive because spend is scaling faster than capability. IDC projects worldwide AI spending to reach $632 billion by 2028, with GenAI at $202 billion (32%).[4] IDC also finds average reported GenAI ROI of 3.7x per dollar invested, but leaders report 10.3x.[6]

The binding constraints are organizational, not algorithmic. In NewVantage's benchmark survey, 92% of Fortune 1000 data and AI leaders cite people and organizational change as the main barrier, and 8% cite technology.[5] Adoption is contested: 31% of employees report sabotaging their company's AI strategy and 94% of C-suite leaders are dissatisfied with their AI solutions.[7] The collateral damage is real: 45% report burnout in transformation programs, and IDC expects the IT skills shortage to hit nine out of ten organizations by 2026, with $5.5 trillion in losses.[12][13]

Top performers build AI as a product capability, underpinned by platforms and unit economics. DORA reports more than 90% adoption of internal developer platforms and 76% adoption of platform engineering teams.[15] They instrument ROI: McKinsey finds KPI tracking is the scaling practice most associated with bottom-line impact, yet fewer than one in five companies do it.[9] They prepare for agents: Deloitte reports 11% have agentic AI in production (38% piloting), while MuleSoft reports 88% pursuing agentic transformation.[17][18]

Treat AI execution as industrial work: measure your funnel, harden AI-ready data, assign accountable owners, invest in MLOps and platforms, manage adoption, and buy partners for outcomes, not pilots.

PART I

Quantifying the Execution Gap

Chapter 1: The "95% Failure" Claim Is Directionally Right, and Often Misused

The claim that "95% of enterprise AI fails" survives because it points at a real pain, even when the metric is sloppy. The story is not a single failure rate. It is a funnel with multiple choke points. Confusing those choke points leads executives to fight the wrong battle.

Start with definitions. A program can "fail" because it never leaves proof-of-concept, because it ships but does not scale, or because it scales but cannot be tied to measurable P&L impact. Different research measures different denominators, so the reported percentages look inconsistent, even when they are describing the same system.

The most value-centric framing is MIT Project NANDA's: 95% of organizations are getting zero measurable P&L return from enterprise generative AI, while 5% of integrated pilots extract significant value. The report also estimates $30 billion to $40 billion already invested. It flags research limitations, including interview-derived figures and a "directionally accurate" warning. Read it as a signal about the distribution of value, not as audited accounting.[1] A more delivery-centric framing comes from S&P Global and 451 Research, which measure pipeline attrition: organizations report scrapping 46% of AI proof-of-concepts before production, and the share abandoning the majority of initiatives rose from 17% to 42% year over year.[2] Gartner's 2026 analysis adds a post-proof-of-concept view: more than half of generative AI projects were abandoned after proof-of-concept by the end of 2025.[3]

BCG's 2025 analysis reaches the same conclusion in a different form: 5% of companies report achieving AI value at scale, while 60% report little to no material value. It is not a project-level failure rate, but it reinforces the same distributional reality.[10]

These are different numbers because they answer different questions. The executive question that matters is which choke point is yours. The table below reconciles the major definitions that drive the public debate.

FAILURE DEFINITION	UNIT OF ANALYSIS	RATE	IMPLICATION	SOURCE
No measurable P&L return from enterprise GenAI	Organization-level value outcome	95%	Activity exists, but value capture is missing at the financial statement level	MIT Project NANDA (2025)[1]
Integrated pilots extracting significant value	Pilot-level value outcome	5%	A small cohort is translating pilots into measurable gains	MIT Project NANDA (2025)[1]
Proof-of-concepts scrapped before production	Project pipeline attrition	46%	Nearly half of PoCs die before broad adoption	S&P Global / 451 Research (2025) [2]
Companies abandoning the majority of AI initiatives	Organization-level portfolio abandonment	42% (up from 17%)	Failure is systemic portfolio churn for many firms	S&P Global / 451 Research (2025) [2]
GenAI projects abandoned after PoC	Project-level abandonment	>50%	Even after clearing feasibility, many programs collapse before scale	Gartner (2026)[3]

The implication is uncomfortable. The center of gravity is not a model problem. It is an execution problem that shows up as value shortfall, pipeline attrition, and post-pilot collapse. With definitions clarified, the next question is obvious: why is spending accelerating while outcomes remain so uneven.

Chapter 2: Follow the Money, Then Follow the Value

Enterprise AI is already a macro budget line. The strategic mistake is to treat it like a lab expense. The budget trajectory says it is becoming part of core operating expenditure, and that makes the execution gap expensive.

IDC projects worldwide AI spending to exceed $630 billion by 2028, with generative AI reaching $202 billion, about 32% of the total.[4][32] IDC's FutureScape framing also places 2025 AI spending at $227 billion and estimates that 67% of that spend comes from enterprises embedding AI into core operations.[31] The scale matters because it sets expectations. When AI becomes a nine-figure line item for large enterprises, "interesting pilots" stop being a defensible output. Investors and boards will demand evidence of cost reduction, revenue lift, or risk reduction.

The ROI data shows a sharp bifurcation. In IDC's 2024 global survey of more than 4,000 business leaders and AI decision-makers, the average reported ROI for generative AI is 3.7x per dollar invested. Leaders report 10.3x. The study is survey-based and self-reported, so it is not a substitute for audited returns, but the spread is still meaningful: the same underlying technology can yield radically different outcomes depending on execution.[6]

Two high-profile examples show what "measurable" looks like when the operating system is in place. JPMorgan's CEO has said the bank spends about $2 billion a year on AI and saves roughly the same amount, framing AI as both a cost center and a productivity lever. It is a CEO-level assertion, not a line-item disclosure, but it signals that large-scale internal value capture is feasible when AI is integrated into operations.[19] Walmart has reported using multiple large language models to create or improve more than 850 million catalog data elements, arguing the work would have required roughly 100 times the headcount without generative AI. That is a management claim rather than a disclosed methodology, yet it captures the right unit of value: throughput in a core operational workflow.[20]

If AI spending is rising and ROI is plausible, why does the median enterprise still struggle? The answer is measurement discipline. McKinsey's global survey finds the scaling practice most associated with bottom-line impact is tracking well-defined KPIs for generative AI solutions, but fewer than one in five respondents report doing it. In other words, many organizations are scaling spend without scaling measurement. That is how you get a large budget line with a small confidence interval around the value.[9]

Chapter 3: The Hidden Cost of Failed AI Is Organizational Debt

Most AI postmortems stop at sunk cost. That is lazy accounting. The larger damage is what repeated failure does to trust, talent, and the willingness to fund the next wave.

The first hidden cost is human. In a global survey of large enterprises, 45% of employees report burnout tied to transformation efforts, and 36% say they would consider quitting. More than half say AI initiatives accelerate fatigue. Those are not "soft" metrics. They are leading indicators of productivity loss and attrition risk in precisely the roles required to scale AI.[12] The second hidden cost is capability erosion. IDC estimates the IT skills shortage will affect nine out of ten organizations by 2026, with $5.5 trillion in losses from delays, impaired competitiveness, and lost business. AI programs that churn in pilot purgatory consume the same scarce talent needed to build the fundamentals. They increase the shortage by wasting the supply.[13]

Trust also has a measurable financial signature. A FTSE Russell analysis of the Fortune 100 Best Companies to Work For finds 3.5x stock market outperformance over 27 years, suggesting that sustained employee trust is not a cultural luxury. It is an economic asset.[27]

The third cost is reputational and regulatory. When AI outputs are wrong in a public context, the remediation cost is often small relative to the trust damage, but it is still quantifiable. In October 2025, Deloitte Australia refunded the final installment of a government contract worth A$440,000, about US$290,000, after a report contained AI-generated errors. The amount is modest. The signal is not. It is a concrete illustration of how governance and quality control failures can become contractual penalties.[23]

Finally, high-profile failures set back internal adoption. Zillow's market value fell by about $30.5 billion from its February 2021 peak to early November 2021 after its iBuying initiative unraveled, a reminder that AI embedded in asset-heavy operations can crystallize into financial losses quickly. The causes were not just algorithmic; they included operational capacity constraints. That pattern repeats in enterprise contexts, where workflow integration and operational limits matter as much as model accuracy.[30]

The execution gap therefore compounds. Each failed project raises the hurdle for the next one. It also drains the very talent needed to close the gap.

PART II

Why Enterprise AI Fails

Chapter 4: Most AI Failures Are Operating Model Failures

The evidence is now overwhelming: technology is rarely the binding constraint. People, incentives, and decision rights are.

NewVantage Partners' 2025 benchmark survey makes the point with an unusually clean split. Ninety-two percent of Fortune 1000 data and AI leaders say the primary barrier to becoming data and AI driven is people and organizational change. Only 8% cite technology.[5] This is not a sentimental argument about culture. It is a quantitative claim about where programs break.

The same theme appears in scaling readiness. In the World Economic Forum and Accenture study on moving beyond experimentation, 74% of companies report challenges adopting AI at scale, and only 16% say they are prepared for AI-enabled reinvention. That gap is a governance and operating model problem. It implies most firms are trying to run an industrial technology through a non-industrial organizational structure.[11]

Worker-level resistance is no longer anecdotal. In the Writer and Workplace Intelligence survey, 31% of employees say they are sabotaging their company's AI strategy, with higher rates among younger workers. The definition includes behavior such as using non-approved tools, refusing training, or refusing to use AI outputs. Regardless of semantics, it quantifies a core reality: adoption is contested inside the firm.[7] The C-suite feels the same strain from the top. In that same research, 94% of executives report dissatisfaction with their organization's AI solutions.[7]

These numbers lead to a sharp conclusion. In most enterprises, AI is treated as a technology program owned by IT, while its value sits in business workflows owned by the line. That is how you get silos, shadow tools, and stalled pilots. Fixing it requires strategic clarity and capability building, not just better prompts. The next chapter examines the mistakes made at the moment programs are initiated.

Chapter 5: The Strategy and Skills Mistakes Are Repeating, and Avoidable

Enterprises have learned to buy cloud. Many have not learned to run AI. The gap shows up at the start, when use cases are chosen and teams are built.

A common strategic failure is to pick use cases that are easy to demo but hard to operationalize. Gartner's data-ready warning captures the mechanism. Gartner forecasts that through 2026, organizations will abandon 60% of AI projects that are not supported by AI-ready data. It is a conditional forecast, not a blanket failure rate, which is precisely why it is useful. It isolates a prerequisite that many programs skip.[8]

Execution breakdown is also a leadership problem, not just a data problem. RAND's analysis, based on interviews with 65 experienced practitioners, highlights business leadership misunderstanding as a common root cause of AI project failure, with data limitations close behind. That sequencing matters because it changes what you fix first.[26]

The skills problem then amplifies the strategic problem. The World Economic Forum estimates that six in ten workers will require training before 2027, while only half have access to adequate training today.[14] That implies a training deficit measured in years, not quarters. Meanwhile, IDC expects the IT skills shortage to hit nine out of ten organizations by 2026, with $5.5 trillion in losses.[13]

The labor market is already signaling where the real demand sits. Job postings for forward-deployed engineers, a role built around embedding technical capability inside client workflows, rose by more than 800% between January and September 2025.[24] That is the opposite of the public narrative, which treats AI talent as a contest for model researchers.

Chapter 6: The Technical Tax Turns Promising Pilots Into Fragile Systems

A prototype is a demonstration. A production system is a liability. The difference is the technical tax: data, integration, monitoring, security, latency, and operations.

Data is the most consistent technical choke point because it is a prerequisite for everything else. Gartner's forecast that 60% of AI projects will be abandoned if unsupported by AI-ready data is a direct statement of this gating effect.[8] Informatica's survey-based research adds operational texture: 78% of IT and data teams report struggling with data orchestration and tooling complexity, along with managing data variety, volume, and quality. It also reports pipeline build times that can reach up to 12 weeks, which is long enough to make many pilots obsolete before they reach production review.[34]

Even when data is available, the production burden is often underestimated. The MLOps literature points to a simple failure mode: teams ship models, then fail to monitor them. ACM Queue reports that roughly half of machine learning practitioners do not monitor model performance in production, citing a production survey of practitioners. That is not a minor omission. It is a decision to fly blind in a system designed to degrade as the world changes.[35]

Generative AI adds a separate reliability tax. Menlo Ventures attributes 15% of failed enterprise generative AI pilots to technical issues, especially around hallucinations. The number is important because it reframes the debate. Hallucinations are not a philosophical problem. They are a failure mode that can kill pilots when governance and grounding are not engineered into the workflow.[36]

The firms that cross this chasm treat infrastructure as strategy. The DORA AI Capabilities Model report finds more than 90% of organizations have adopted internal developer platforms, and 76% have dedicated platform engineering teams. Those numbers are a proxy for a broader shift: enterprises are rebuilding the software delivery substrate to make AI deployable, observable, and governable.[15]

PART III

What the Top 5% Do Differently

Chapter 7: The Winners Build an AI Operating System, Not a Collection of Models

The top performers do not "do AI projects." They build an AI capability that ships repeatedly. That requires a stable operating system: governance, platforms, teams, and decision rights.

Platform engineering has become the backbone of this operating system. The DORA data is revealing because it signals how quickly the market is converging. More than 90% of organizations report adopting internal developer platforms, and 76% report dedicated platform engineering teams.[15] In mature software organizations, this kind of platform investment is not cosmetic. It is how you standardize deployment, security controls, monitoring, and developer experience.

Governance structures are also evolving, but unevenly. A classic pattern is the AI Center of Excellence. Harvard Business Review reported that in a survey of U.S. executives from large firms using AI, 37% said they had already established an AI Center of Excellence. The data is dated, but the signal still matters: even among large firms that were early AI adopters, the majority had not institutionalized governance in a dedicated structure.[28]

The winners fill that vacuum with clear accountability. JPMorgan's approach illustrates the pattern. Its CEO frames AI as a multi-billion-dollar internal capability, not a product experiment.[19] At the operating level, Walmart illustrates the same point in a different industry. When a firm reports improving more than 850 million catalog data elements with generative AI, it is describing a workflow embedded in core operations, not a demo.[20]

Chapter 8: The Best Firms Instrument Value Like Reliability

AI ROI cannot be managed if it cannot be observed. The firms that reach measurable ROI treat value as something to instrument, not something to narrate in quarterly decks.

The gap is measurable. In McKinsey's global survey, tracking well-defined KPIs for generative AI solutions is the scaling practice most associated with bottom-line impact, yet fewer than one in five respondents report doing it.[9] IDC's work adds an even sharper contrast: average reported generative AI ROI is 3.7x per dollar invested, while leaders report 10.3x. Even if self-reported, the spread implies that measurement and operating discipline are creating order-of-magnitude differences in value capture.[6]

A second signal comes from the vendor-led pilot economy. In an IDC survey reported by CIO.com, nearly two-thirds of CIOs said vendor-led AI proof-of-concepts have a 90% failure rate. The same reporting describes organizations running dozens of pilots, with only a handful reaching production.[16]

The winners therefore connect MLOps and business measurement. They treat cost and value as a single pipeline, with unit economics that can be tracked daily. BT Group reported cutting AI deployment time from about six months to six days through its internal AI Accelerator platform.[21] ClearML's practitioner survey suggests the market is moving in the same direction: 85% of respondents reported having a dedicated MLOps budget in 2022, with another 14% planning to add one in 2023.[22]

Chapter 9: What "AI at Scale" Actually Looks Like

The most useful case studies share an unglamorous trait. They are deeply integrated into core workflows, and they are governed like production systems.

JPMorgan provides the archetype for internal value capture. Its CEO has publicly framed AI as a $2 billion annual spend that saves roughly $2 billion, treating AI as a managed productivity investment.[19] Whether the exact accounting mapping is disclosed or not, the intent matters: the organization is willing to invest at scale because it believes it can measure savings at scale.

Walmart provides the archetype for operational throughput. In its earnings-call disclosures, the company reported using multiple large language models to create or improve more than 850 million catalog data elements, arguing the same effort would have required about 100 times the headcount without generative AI.[20]

BT Group provides the archetype for platform leverage. It reported reducing AI deployment time from six months to six days through an internal AI Accelerator. This is the kind of improvement that changes the economics of experimentation. When deployment is slow, every pilot is expensive and politically fragile. When deployment is fast, pilots become a portfolio and scaling becomes a choice rather than a miracle.[21]

Vendor case studies can be illustrative if treated carefully. Dataiku reports a client building more than 150 AI products with over 750 AI creators and cutting time-to-market by 85%. It is a marketing case study rather than an independent benchmark, so it should be read as an existence proof, not a market average.[29]

These cases converge on a repeatable pattern. The winners pick workflows with clear unit economics, invest early in data and integration, ship through reusable platforms, and treat adoption as a managed change program.

PART IV

The Implementation Market and the Road to 2028

Chapter 10: Buyers Are Frustrated, and the Data Explains Why

Enterprise buyers are not asking for inspiration. They are asking for outcomes. The frustration is measurable, and it is widening.

The most direct signal is dissatisfaction. In the Writer and Workplace Intelligence study, 94% of C-suite leaders report dissatisfaction with their organization's AI solutions, and 98% say they want their AI vendors to help set their AI strategy and vision.[7] That combination is revealing. Buyers are delegating vision externally because internal capability is weak, then they are unhappy with what they receive.

BUYER FRUSTRATION SIGNAL	MEASUREMENT	FIGURE	SOURCE
C-suite dissatisfaction with AI solutions	Executive sentiment	94% dissatisfied	Writer / Workplace Intelligence (2025)[7]
Desire for vendors to set AI vision	Executive expectations	98%	Writer / Workplace Intelligence (2025)[7]
Employee resistance to AI program	Self-reported sabotage	31%	Writer / Workplace Intelligence (2025)[7]
Vendor-led PoC failure rate	CIO-reported pilot outcomes	90%	IDC via CIO.com (2024)[16]
PoCs scrapped before production	Internal delivery conversion	46%	S&P Global / 451 Research (2025)[2]
Firms abandoning most initiatives	Portfolio-level churn	42% (up from 17%)	S&P Global / 451 Research (2025)[2]
GenAI projects abandoned after PoC	Post-pilot collapse	>50%	Gartner (2026)[3]

The procurement implication is straightforward. The market will reprice toward outcome accountability. The vendors that win will be those that can take responsibility for integration, controls, and measurable value, not just model selection.

Chapter 11: The Rise of the AI Operator, and Why Consulting Must Change

Traditional consulting is built for advice. Enterprise AI requires operations. The gap between those two is now a market opportunity.

Labor market signals show the shift in real time. Job postings for forward-deployed engineers rose by more than 800% between January and September 2025. This role is designed to sit at the boundary between technical systems and client workflows, shipping production outcomes rather than recommendations.[24]

Revenue signals point in the same direction. CB Insights estimates that private AI agent solutions, plus the underlying LLM layer, generated more than $10 billion in revenue in 2024. That is a meaningful pool for a category that barely existed a few years ago.[25]

Forrester's view of the market aligns with that shift. In its 2026 predictions, it frames AI as moving from hype into operational work, where value depends on execution discipline rather than experimentation. As AI platforms automate more delivery tasks, service providers will be pushed toward repricing and shared-risk models.[33]

The emerging segmentation can be described without jargon. There are firms that sell advice, firms that sell implementation labor, firms that sell platforms, and firms that bundle platforms with embedded operators who build and run workflows inside the client environment. The last category is still small, but it is growing because it matches the failure modes quantified in Chapter 10.

Chapter 12: Agentic AI Will Widen the Gap Through 2028

Agentic AI changes the execution equation because it targets entire workflows, not isolated tasks. It also changes the risk profile because it can take actions, not just generate text. This will widen the gap between enterprises that can run AI in production and those that cannot.

The adoption signal is already split between intent and reality. Deloitte's Tech Trends 2026 reports that only 11% of organizations have agentic AI systems in production, even though 38% are piloting them.[17] That is a classic execution gap, now expressed in the agent category. At the same time, Salesforce and MuleSoft report that 88% of organizations are pursuing partial or full agentic transformation, a number that reflects ambition more than readiness.[18]

Spend forecasts suggest the agent wave will not be a niche. IDC expects generative AI spending to reach $202 billion by 2028, about 32% of total AI spending.[4] As agentic tools become embedded in enterprise software, a meaningful share of that spend will migrate from "chat interfaces" into workflow automation, decision support, and orchestration.

Three factors determine whether agentic AI narrows or widens the execution gap. Integration is first. Agents are only as good as the systems they can act on. Governance is second. If an organization cannot define decision rights for an AI assistant, it will fail faster for an agent that can trigger actions. Observability is third. Without monitoring, an agentic workflow can degrade quietly, then fail loudly in a customer-facing context.

The conclusion is sharp. Through 2028, the winners will be the firms that treat agentic AI as a production discipline from day one. The laggards will run agent pilots as theater, then discover that operational risk scales faster than the demo.

Recommendations

1. Replace the "AI success" narrative with a measurable funnel

Enterprises should stop debating a single failure rate and start measuring their own funnel. A practical funnel has at least three stages: proof-of-concept conversion, production scaling, and measurable business impact. External benchmarks show why this matters. Organizations report scrapping 46% of AI proof-of-concepts before production, and more than half of generative AI projects are abandoned after proof-of-concept.[2][3] If you do not track your own conversion rates and their causes, you will misdiagnose where the program is breaking.

2. Treat AI-ready data and integration as a prerequisite, not a workstream

The most predictive technical risk factor is data readiness. Gartner forecasts that 60% of AI projects unsupported by AI-ready data will be abandoned through 2026.[8] Informatica reports that 78% of IT and data teams struggle with data orchestration and tooling complexity, and pipeline build times can reach up to 12 weeks.[34] Leaders should fund a dedicated AI-ready data program with clear deliverables. This is the cheapest place to buy down downstream failure risk.

3. Put one accountable owner over workflow, data, and controls

Most enterprises split accountability across IT, data teams, and business functions. That structure makes it almost impossible to move from pilot to scaled workflow change. The organizational evidence is decisive: 92% of Fortune 1000 data and AI leaders cite people and organizational change as the primary barrier, not technology.[5] Assign a single accountable owner for each AI product in production.

4. Invest in MLOps and platform engineering as the delivery substrate

Platform capability is the quiet differentiator. More than 90% of organizations report adopting internal developer platforms, and 76% have dedicated platform engineering teams.[15] ClearML's survey suggests MLOps is becoming a budgeted line item, with 85% reporting a dedicated budget.[22] Standardize deployment, monitoring, access controls, and rollback. AI will not scale on bespoke pipelines.

5. Run adoption as change management with measurable behaviors

AI programs fail when adoption is treated as organic. Thirty-one percent of employees report sabotaging their company's AI strategy, and 45% report burnout tied to transformation fatigue.[7][12] Design adoption like a product rollout: targeted training, opt-in onboarding, clear policy on approved tools, and measurable usage metrics tied to workflows.

6. Buy AI implementation based on outcome accountability, not reputation

Nearly two-thirds of CIOs report a 90% failure rate for vendor-led AI proof-of-concepts, and 94% of the C-suite reports dissatisfaction with AI solutions.[16][7] Procurement should shift selection criteria: require production references with measurable outcomes, demand transparency on integration and governance responsibilities, and price contracts around milestones that reflect real value capture.

Methodology Note

This report was produced using 900 Labs' adversarial multi-model research pipeline. The pipeline is designed to reduce the two failure modes that plague AI research synthesis: citation hallucination and selective evidence. It combines structured source retrieval across primary research, analyst reports, and enterprise disclosures, with cross-model adversarial review where independent models challenge each other's claims, definitions, and denominators. Claims that survive cross-model scrutiny are then subjected to primary-source verification, prioritizing original PDFs, official transcripts, and first-party disclosures over secondary summaries. Where sources use survey-based or self-reported metrics, the report treats them as directional signals rather than audited financial facts, and it labels the unit of analysis and the denominator explicitly to avoid false precision.

References

[1] MIT Project NANDA, "The GenAI Divide: State of AI in Business 2025" (July 2025). Archive

[2] S&P Global / 451 Research, "AI experiences rapid adoption but with mixed outcomes" (2025). Link

[3] Gartner, "Why Half of GenAI Projects Fail: Avoid These 5 Common Mistakes" (2026). Link

[4] IDC, "AI spending to reach $632 billion by 2028" (Aug 2024). Link

[5] NewVantage Partners, "2025 AI & Data Leadership Executive Benchmark Survey" (Dec 2024). Link

[6] IDC, "2024 Business Opportunity of AI: GenAI Delivering New Business Value and Increasing ROI" (2024). Link

[7] Writer / Workplace Intelligence, "2025 Enterprise AI Adoption Report" (2025). Link

[8] Gartner, "Lack of AI-Ready Data Puts AI Projects at Risk" (Feb 2025). Link

[9] McKinsey & Company, "The state of AI in 2025: Agents, innovation, and transformation" (2025). Link

[10] BCG, "The Widening AI Value Gap: Build for the Future" (Sept 2025). Link

[11] WEF / Accenture, "AI in Action: Beyond Experimentation to Transform Industry" (2025). Link

[12] Emergn, "The Global Intelligent Delusion" (Oct 2025). Link

[13] IDC via BusinessWire, "IT Skills Shortage Expected to Impact Nine out of Ten Organizations by 2026" (May 2024). Link

[14] WEF, "Future of Jobs Report 2023" (Apr 2023). Link

[15] Google Cloud / DORA, "2025 DORA AI Capabilities Model Report" (2025). Link

[16] CIO.com, "CIOs' lack of success metrics dooms many AI projects" (Dec 2024). Link

[17] Deloitte, "Tech Trends 2026" (Feb 2026). Link

[18] Salesforce / MuleSoft, "2026 Connectivity Benchmark Report" (2026). Link

[19] JPMorgan Chase / Bloomberg TV via Investing.com (Oct 2025). Link

[20] Walmart Inc., Q2 FY25 Earnings Call Transcript (Aug 2024). Link

[21] BT Group, "BT Group unveils AI Accelerator" (Oct 2022). Link

[22] ClearML, "MLOps in 2023: What Does the Future Hold?" (Jan 2023). Link

[23] Fortune, "Deloitte caught using AI in $290K report" (Oct 2025). Link

[24] Fast Company, "Postings for this AI job are up 800%" (Nov 2025). Link

[25] CB Insights, "Future of Professional Services" (2025). Link

[26] RAND Corporation, "The Root Causes of Failure for AI Projects" (2024). Link

[27] Great Place to Work / FTSE Russell (2025). Link

[28] HBR, "How to Set Up an AI Center of Excellence" (Jan 2019). Link

[29] Dataiku, "Making Enterprise AI an Organizational Asset" (2025). Link

[30] Bloomberg, "Zillow's Drop Brings Market Value Loss to $30 Billion From Peak" (Nov 2021). Link

[31] IDC via BusinessWire, "IDC Unveils 2025 FutureScapes" (Oct 2024). Link

[32] IDC, "A Deep Dive Into IDC's Global AI and GenAI Spending" (2024). Link

[33] Forrester, "Predictions 2026: AI Moves From Hype To Hard Hat Work" (Oct 2025). Link

[34] Informatica, "Data Integration Tools: Challenges and Key Trends 2024" (2024). Link

[35] ACM Queue, "Unsolved Problems in MLOps" (2025). Link

[36] Menlo Ventures, "2024: The State of Generative AI in the Enterprise" (2024). Link

900 LABS

Superintelligence, Operationalized

900labs.com