

100 Words a Minute. 20 of Them Right.

Companies are measuring AI progress by how fast they adopt it. The groundwork that makes AI accurate — the data, the governance, and the people who do that work — is the part nobody is funding.

Daniel Rodriguez • Founder & CEO, Infragil • September 2026

I was sitting in a class recently, listening to a room full of senior people describe their AI programs, and the same shape kept appearing in every story. Different industries, different maturity levels, same shape.

Everyone was moving fast. Almost nobody could tell you whether they were moving in the right direction.

It reminded me of a typing test.

If you type 100 words a minute and only 20 of them are right, you do not type 100 words a minute. You produce errors at speed.

That is where a lot of AI programs are. The adoption metrics look tremendous. Usage is up. Pilots are everywhere. Board decks are full of green. And underneath it sits a data estate nobody has cleaned, a governance model nobody owns, and a set of people who were never given the time or the training to do the part that actually makes the output correct.

The usual caveat, because it matters: I run an AI-native company, I use these tools every day, and they remain the biggest force multiplier I have seen in three decades. This is not an argument against AI. It is an argument about what we are counting.

The statistic I almost published

I started this article with a number in my head. I was fairly sure Gartner had reported that only 20% of AI projects are successful. I had heard it, it matched what I was seeing, and it would have made a clean opening line.

So I went to check it. Gartner never said it.

There is no Gartner press release, prediction or survey with that finding. What Gartner actually published, in April 2026, is that 28% of AI use cases in infrastructure and operations fully succeed and meet ROI expectations, and 20% fail outright. Somewhere between that release and my memory, "20% fail" had become "20% succeed."

Once I started pulling threads, it got worse. Four of the most-quoted AI failure statistics in circulation do not survive contact with their sources:

- **"RAND found 80% of AI projects fail."** RAND did not find that. RAND's 2024 report repeats the figure in passing as "by some estimates," and the footnote attributes it to a Fortune article from

July 2022. RAND's actual research is 65 practitioner interviews producing five root causes — no failure rate at all.

- **"MIT found 95% of GenAI pilots fail."** The July 2025 MIT NANDA paper says 95% of organizations are getting zero return — a different unit and a different measure. Wharton's Kevin Werbach has said publicly he read the document repeatedly and could not locate where the number comes from. It is not peer reviewed, and the authors commercialize the category of solution it recommends.
- **"Gartner says 85% of AI projects fail."** Gartner's 2018 prediction was that 85% of AI projects would deliver erroneous outcomes due to bias in data, algorithms or teams. That is a statement about error, not project failure, and its window closed in 2022.
- **"87% of data science projects never reach production."** This traces to a 2019 VentureBeat piece covering a conference panel. It cites no study. It is a remark someone made on stage, and it is routinely attributed to Gartner.

I want to be careful not to be smug about this, because I was about to publish one of them myself.

But look at what that pile of numbers is. It is an entire discourse about AI accuracy, built out of statistics nobody checked, repeated at speed, by people who were confident. That is the typing test. That is the whole article in one paragraph, and I found it by accident while trying to write the opening.

What the defensible numbers actually say

Here is what survives verification. Note how much more modest — and more useful — it is.

Finding	Figure
AI use cases in infrastructure & operations that fully succeed and meet ROI expectations (Gartner, Apr 2026, n=782)	28%
... that fail outright, same survey	20%
Organizations that have successfully scaled AI across multiple business units (Gartner, Sep 2026, n=1,303)	22%
Organizations that have moved 40% or more of AI pilots into production (Deloitte, Jan 2026, n=3,235)	25%
Respondents attributing at least some EBIT impact to AI (McKinsey, Aug 2026, n=1,719)	37%
Digital trust professionals reporting AI ROI met or exceeded expectations (ISACA, May 2026, n=3,400+)	22%

No single credible "AI success rate" exists, because every source measures a different unit. But the range is consistent, and it is not flattering: somewhere between a fifth and a third of this work is paying off, depending on what you count.

And here is the part I have to be honest about, because it cuts against the argument I want to make. **Data quality is not the number one cause of AI project failure.** In that same Gartner survey, the top cited cause was "expecting too much, too fast" at 57%. Poor data quality or limited availability came second, at 38%. RAND's interviews put leadership and communication failures first and data second. McKinsey points at organizational absorption capacity and workflow redesign.

So the honest version of my thesis is narrower than the one I started with, and I think it is stronger. Data readiness is the most persistent technical cause of failure, and it is the only one with a documented positive correlation to success: Gartner found in April 2026 that organizations with successful AI initiatives invest up to four times more, as a percentage of revenue, in their data and analytics foundations than poor performers.

Expecting too much too fast is the leading cause of failure. Unready data is the thing that makes the expectation unrealistic in the first place.

Everybody says their data is ready

This is the finding that surprised me most, and it reframed the whole piece.

I expected the research to show that companies know their data is not AI-ready and are pressing ahead anyway. The opposite is true. Companies are confident. It is the same companies' own answers, inside the same surveys, that give them away.

What they say	What they also say
85% have a clear data strategy	Only 18% say their data is fully governed
84% are confident in their data's accuracy and completeness	About 80% say AI is constrained by limited data access
87% perceive their data as AI-ready	43% name data readiness as the most significant barrier
71% say they connect AI well to business goals	Only 31% have metrics actually tied to KPIs
74% plan to deploy AI agents within two years	Only 21% have a mature governance model for them

The first four rows are two surveys: Cloudera's Data Readiness Index, fielded January to March 2026 across 1,270 IT leaders, and the fourth annual Precisely study with Drexel University's LeBow College of Business, 505 data and analytics leaders. The last row is Deloitte, 3,235 leaders across 24 countries.

Two of those three are vendor-commissioned, and I will say so plainly: Cloudera sells data platforms and Precisely sells data integrity tooling. Both are measuring a problem their products exist to solve. The Precisely study has a named academic co-author at Drexel and a jointly developed instrument, which is better than most vendor research, but it is not independent. Deloitte sells AI consulting. Weigh them accordingly.

What makes the pattern credible is not any one study. It is that the contradiction shows up inside each of them, between answers given by the same respondents. Eighty-five percent have a strategy and eighteen percent have governance. That gap is not vendor spin. That is people describing an intention and a reality in the same breath and not noticing.

Gartner, independently, found that 63% of organizations either do not have the right data management practices for AI or are not sure whether they do. "Or are not sure" is doing a lot of work in that sentence.

You cannot point AI at a data lake and get the truth

There is a belief, and I hear it from senior people constantly, that the data problem is solved by collection. Get it all in one place, point the model at it, and the model will sort it out.

That is not how retrieval works.

The best technical treatment I found is a peer-reviewed paper from Deakin University's Applied Artificial Intelligence Institute, presented at CAIN 2024, which catalogues seven distinct failure points in retrieval-augmented systems from three real case studies. The failures are not exotic. The answer is missing from the corpus. The answer exists but does not rank into the retrieved set. The answer is retrieved but does not survive consolidation into context. And my favourite, because it is the one that maps directly onto enterprise data:

The answer is in the context and the model still fails to extract it — "typically due to excessive noise or contradicting information." — Barnett et al., Deakin University, CAIN 2024

Contradicting information. That is every enterprise I have ever worked in. Three systems of record with three different customer addresses. A policy document superseded twice but never archived. Two reports with the same title and different numbers, and the only person who knew which one was authoritative left in 2023.

A human analyst handles that with judgment and institutional memory. A retrieval system handles it by picking one, confidently.

Then there is the property that makes people most uncomfortable when I explain it: the same question can produce a different answer. That is not a bug in a particular product. When a model is served at scale, the batch it happens to be processed in varies with load, and that changes the arithmetic. As the team at Thinking Machines put it, from the perspective of an individual user, the other concurrent users are not an input to the system but a nondeterministic property of it.

Sit with the governance implication of that for a moment. You have a control that cannot be tested by re-running it and comparing the output. Every audit methodology I have used in 29 years assumes reproducibility. That assumption no longer holds, and almost nobody has updated their control testing to account for it.

And the retrieval layer is a security boundary, not just a quality one. Microsoft's own deployment guidance for Copilot says it uses data the user already has permission to access — which means it inherits whatever permission mess already exists, and enforces nothing about whether those permissions were ever correct. Microsoft then instructs administrators to go find sites and files that are

overshared, ownerless, or inactive, and to repair broken permission inheritance. Read the framing of that guidance closely: it is written for organizations preparing for Copilot, or making adjustments after Copilot is enabled.

The vendor's own documentation contemplates that you turned it on first and will fix the permissions afterward. That is not a criticism of the vendor. They are describing their customers accurately.

The work nobody wants to fund

Everything above is fixable. It is fixable by people doing unglamorous, well-understood work: cataloguing data, assigning ownership, fixing lineage, writing retention rules, cleaning duplicates, repairing permissions, and testing outputs. None of it demos well.

So it does not get funded.

- **dbt Labs' 2026 survey of analytics engineering teams found 57% reporting increased warehouse and compute spending — and only 36% reporting increased team budgets.** More money for the machines. Less for the people operating them.
- **The 2026 IANS and Artico Search budget research found security budgets grew 5% on average — but the median growth was zero.** Sixty-four percent of CISOs asked for an increase and 45% got one. The "AI money" is real, and it is concentrated in a handful of organizations while the typical security leader got nothing.
- **The Conference Board found in July 2026 that 55% of workers use generative AI or AI agents daily or weekly, while only 33% had used employer-provided AI training in the past six months — and 28% say their organization provides no AI training at all.** We gave them the tools. We did not teach them the job.
- **Proofpoint's 2026 survey of 1,600 CISOs found 78% are expected to manage AI-related risks without a proportional increase in resources or expertise.** Not "under-resourced" in the abstract. Specifically: new risk, same budget, same headcount.

BCG surveyed 11,749 workers across 14 markets this June and found the thing that should end the argument. Having a clear AI strategy lifted measurable business impact by 25 percentage points. Better tools alone lifted it by five.

Strategy is worth five times what tooling is worth, and tooling is where the money went.

The gap between the three groups is as wide as the Grand Canyon

Here is the part I have been circling for a while, and the part that I think is least well understood.

There are three groups who have to agree for any of this to work, and they do not speak the same language.

Executives do not understand the intersection.

Grant Thornton surveyed 950 business leaders in early 2026 and found 48% of boards had not set AI governance expectations and 46% had not integrated AI risk into ongoing oversight — while 75% had approved major AI investments. Funding without governing. Deloitte's boardroom research found 66% of boards still have limited to no knowledge or experience with AI. ISACA found only 38% of practitioners are confident in their board's understanding of AI risk.

And the executives do not agree with each other either. In that same Grant Thornton data, 54% of COOs were concerned about agentic AI compliance — against 20% of CIOs and CTOs. On whether the workforce is ready for AI, 29% of CIOs and CTOs said yes; 6% of COOs agreed. That is a 23-point gap between peers sitting in the same meetings.

Security leaders do not fully understand the data requirements.

I include myself in this, and I have spent my career in it. My profession is trained to protect data, classify it, control access to it and prove the controls work. We are not, as a discipline, trained in data quality, lineage, schema drift, or what makes a dataset representative enough to train on. Those are different skills, and most of us do not have them.

Logicalis surveyed over 1,000 CIOs and found 62% admit to compromising on governance because of limited knowledge, only 44% say they fully grasp the risks of AI adoption, and 89% describe their approach as learning as we go. Okta found only 31% of CISOs feel fully aligned with their C-suite and board on acceptable AI risk levels.

Data leaders do not fully understand the governance requirements.

And the reverse holds. Excellent data engineers routinely build pipelines with no thought to retention obligations, cross-border transfer restrictions, or the fact that a training set assembled from production data has just created a copy of regulated information in a place no one has inventoried.

dbt Labs found 41% of analytics teams citing ambiguous data ownership as a core challenge. Not disputed ownership. Ambiguous — nobody knows whose job it is.

The most striking number I found in all of this research is from the 2026 AI and Data Leadership Executive Benchmark Survey, now in its fifteenth year: 93.2% of senior data executives identified culture and change management as the greatest impediment to becoming data-driven. Only 6.8% blamed technology. The sample is small and self-selected — about 110 executives, invitation-only, skewed toward large and engaged organizations — so treat it as directional. But the direction has been the same every year for fifteen years.

Fifteen years of asking senior data leaders what is in their way, and the answer has never once been the technology.

The part that looks like the last twenty years of my career

There is a human cost to this, and it is the part of the pattern I recognize fastest, because I have lived it.

The security profession has spent two decades being handed responsibility without authority, scope without headcount, and expectations set by people who did not understand the work. The measured results are what you would expect. ISC2's 2025 workforce study, with just over 16,000 respondents, found 48% exhausted from trying to keep current and 47% overwhelmed by workload expectations, with only 34% reporting appropriate staffing. Splunk found nearly two-thirds of security teams experiencing moderate to significant burnout. And in Splunk's prior-year research, 29% of CISOs said they received adequate budget — while 41% of board members believed the budgets were adequate. The gap is not just in resources. It is in perception, and it sits at the top.

Now I want to be careful, because here I am going past the evidence. My strong impression, from conversations and from that classroom, is that data and AI practitioners are now absorbing the same pattern. I could not find current research that establishes it. There is no 2026 equivalent of the ISC2 study for data teams. The closest direct evidence is a 2021 survey of 600 data engineers, commissioned by two vendors selling the remedy, which found 91% frequently receiving requests with unrealistic or unreasonable expectations. That figure is five years old, it predates generative AI entirely, and its sponsorship is exactly the kind of thing I flagged earlier in this article. I am not going to lean on it.

What I can point to is adjacent and current. Dice found burnout among tech professionals rising from 31% to 46% in a year. Stack Overflow, surveying more than 49,000 developers, found trust in AI accuracy falling to 29% from 40%, with 66% reporting they spend more time fixing almost-right AI-generated code than they would have spent writing it. BCG found 47% of workers spending more time managing AI than doing their actual work.

One more honest note, because the picture is not uniformly bleak and I do not want to overstate it. The same ISC2 study found 68% overall job satisfaction and 78% planning to stay in the field. BCG found 67% of regular AI users reporting improved job satisfaction. The strain and the engagement coexist. People mostly like this work. They are being asked to do it in conditions that do not let them do it well, which is a different and in some ways more solvable problem than people hating their jobs.

Treat the data and AI teams the way your organization treated security teams for twenty years, and you will get the same outcome you got with security teams for twenty years.

The playbook

- 1 Measure accuracy, not adoption.** Usage metrics are the words-per-minute score. Pick a small number of outputs your AI systems produce, have a qualified human grade them for correctness on a schedule, and report that number next to the usage number. If you cannot grade the output, you are not ready to scale the system that produces it.
- 2 Fund the data work explicitly, as its own line item.** Gartner found that organizations with successful AI initiatives invest up to four times more in data and analytics foundations. If data preparation, cataloguing and governance are buried inside a project budget, they lose every prioritization fight to the visible feature. Make them a line the CFO can see.
- 3 Fix permissions before you enable retrieval, not after.** Your AI assistant inherits your access control mistakes and makes them instantly discoverable at scale. Run the oversharing remediation — ownerless sites, company-wide sharing links, broken inheritance — as a prerequisite to

deployment, not a follow-up item. Microsoft's own guidance tells you to do this. Almost nobody does it first.

- 4 Put the three groups in the same room, on purpose, repeatedly.** Not a steering committee that meets quarterly to receive a status deck. Working sessions where the security lead has to explain the control requirements to the data team, and the data team has to explain lineage and quality to the security lead, and both have to explain the constraints to the executive who approved the budget. The gap closes through contact, and nothing else closes it.
- 5 Name an owner for AI risk, in writing.** Grant Thornton found 78% of organizations lack confidence they could pass an independent AI governance audit within 90 days. The most common reason I see is not incompetence. It is that three functions each assume one of the other two has it. Write the name down.
- 6 Cross-train deliberately, and budget the hours.** Send your security people to data governance training. Send your data engineers to privacy and records-retention training. Send both to whatever your legal team runs on regulated data. This is a small, cheap, unglamorous intervention with a larger return than most tooling purchases, and it is the single thing I would do first.
- 7 Check your numbers before you repeat them.** Including the ones in this article. Every figure here has a named source and a date in the list below. If a statistic about AI has been travelling without a source attached, assume it is wrong until you have found the original — because four of the most-quoted ones are.

What this one is about

The three pieces I have written before this one each landed on the same kind of point. Experience is the guardrail. Access is not ability. Disclosure is not visibility.

This one is simpler. Speed is not progress.

Every organization I talk to is measuring how fast it is adopting AI. Very few are measuring whether the output is correct, and fewer still are funding the work that would make it correct. The data has to be cleaned by somebody. The governance has to be owned by somebody. The permissions have to be fixed by somebody. Those people exist in your organization right now, they know what needs doing, and in most companies they have been told to do it on top of their existing job, with no additional budget, by a leadership team that believes the problem was solved when the contract was signed.

The typing test does not care how fast you go. It counts the errors.

You cannot point AI at a mess and get the truth. You can only get a confident, fluent, well-formatted version of the mess — faster than before, and harder to question.

At Infragil, we help companies use AI responsibly and effectively — which starts with the unglamorous part. Data readiness assessment, AI governance ownership, permission remediation before deployment, and the cross-training that closes the gap between your security, data and executive teams. If your AI

adoption numbers look great and nobody can tell you whether the output is right, that is the gap we close. Let's talk.

Daniel Rodriguez is the Founder & CEO of Infragil, an AI advisory and cybersecurity practice, and creator of Skopos, an AI-native third-party risk management platform. He has spent 29 years in cybersecurity, including CISO and VP roles.

Sources

- Gartner (Apr 7, 2026), Melanie Freeze: 28% of I&O AI use cases fully succeed and meet ROI expectations; 20% fail outright; 57% cited "expected too much, too fast"; 38% cited poor data quality or limited data availability. Survey of 782 I&O leaders, Nov-Dec 2025. Scope is infrastructure and operations, not all enterprise AI.
- Gartner (Sep 1, 2026), Tina Nunno: only 22% of organizations have successfully scaled AI across multiple business units. Survey of 1,303 functional leaders at organizations with \$50M+ revenue, Jan-Apr 2026.
- Gartner (Feb 26, 2025), Roxane Edjlali: 63% of organizations either do not have or are unsure whether they have the right data management practices for AI. Surveys of 1,203 and 248 data management leaders, 2024.
- Gartner (Apr 16, 2026), Rita Sallam: organizations with successful AI initiatives invest up to four times more, as a percentage of revenue, in data and analytics foundations. Survey of 353 data, analytics and AI leaders, Nov-Dec 2025.
- Gartner (Jun 25, 2025), Anushree Verma: over 40% of agentic AI projects predicted to be canceled by end of 2027. Note this is a prediction, based on a January 2025 poll of 3,412 webinar attendees — a self-selected audience.
- Gartner (2018), Whit Andrews: the widely misquoted "85%" prediction was that 85% of AI projects would deliver erroneous outcomes due to bias in data, algorithms or teams, through 2022 — not that 85% of projects fail.
- RAND Corporation (Aug 13, 2024), Ryseff, De Bruhl & Newberry: five root causes of AI project failure, from 65 semi-structured practitioner interviews conducted Aug-Dec 2023. Leadership and communication ranked first; data second. The "80% of AI projects fail" figure appears only as "by some estimates," footnoted to a Fortune article of July 26, 2022 — it is not a RAND finding.
- MIT NANDA, "The GenAI Divide: State of AI in Business 2025" (July 2025): claims 95% of organizations are getting zero return, based on 52 interviews and 153 survey responses. Not peer reviewed. Criticism of the unlocatable derivation from Kevin Werbach (Wharton), reported by Futurium; further analysis by Robert Wiblin (80,000 Hours) estimating the pilot-level failure rate nearer 75%.
- VentureBeat (Jul 19, 2019): origin of the unsourced "87% of data science projects never reach production" claim — conference panel coverage citing no study.
- Deloitte, "State of AI in the Enterprise" 2026 (Jan 21, 2026): 21% report a mature governance model for autonomous agents while 74% plan agent deployment within two years; 25% have moved 40%+ of pilots into production. 3,235 leaders across 24 countries, fielded Aug-Sep 2025.
- Cloudera, "The Data Readiness Index" (Apr 14, 2026): 85% report a clear data strategy; 84% confident in data accuracy and completeness; ~80% say AI is constrained by limited data access; 71% say most data is governed; 18% say data is fully governed. 1,270 IT leaders, fielded by Researchscape Jan 22-Mar 3, 2026. Vendor-commissioned — Cloudera sells data platforms.
- Precisely and Drexel University LeBow College of Business, "2026 State of Data Integrity and AI Readiness": 87% perceive their data as AI-ready; 43% cite data readiness as the most significant barrier; 71% claim business-goal alignment while 31% have metrics tied to KPIs; 83% have an ongoing data governance program. 505 data and analytics leaders, fielded H2 2025. Vendor-sponsored with academic co-authorship (Prof. Murugan Anandarajan).
- McKinsey, "The State of AI in 2026: On the Road to ROI" (Aug 2026): 37% attribute at least some EBIT impact to AI; 44% report AI scaling across the enterprise. 1,719 participants across 97 nations, fielded May 4-Jun 8, 2026, GDP-weighted.

- Barnett, Kurniawan, Thudumu, Brannelly & Abdelrazek (Applied Artificial Intelligence Institute, Deakin University), "Seven Failure Points When Engineering a Retrieval Augmented Generation System," CAIN 2024 (arXiv:2401.05856): peer-reviewed; failure point 4 attributes extraction failure to "excessive noise or contradicting information."
- Thinking Machines Lab (Sep 10, 2025), "Defeating Nondeterminism in LLM Inference": batch-size variation under load as the primary source of nondeterministic output; concurrent users as "a nondeterministic property of the system."
- Microsoft Learn, Copilot deployment and governance guidance: Copilot uses data the user already has permission to access; administrators are directed to remediate overshared, ownerless and inactive content and repair broken permission inheritance, either before enabling Copilot or after.
- IBM and Ponemon Institute, "Cost of a Data Breach Report 2025": 63% of breached organizations have no AI governance policies. Of organizations that experienced AI-related security incidents, 97% lacked proper AI access controls — note this subset is roughly 78 organizations of 600 surveyed.
- dbt Labs, "2026 State of Analytics Engineering" (Apr 14, 2026): 57% report increased warehouse and compute spending against 36% reporting increased team budgets; 71% cite incorrect or hallucinated outputs reaching stakeholders; 41% cite ambiguous data ownership. 363 respondents.
- IANS Research and Artico Search, "2026 Security Budget Report" (Sep 17, 2026): security budgets grew 5% on average but median growth was 0%; 64% of CISOs requested an increase and 45% received one; 69% named AI their top net-new spending priority. 500+ security executives, fielded Apr–Aug 2026.
- Proofpoint, "2026 Voice of the CISO" (Sep 9, 2026): 78% of CISOs are expected to manage AI-related risks without a proportional increase in resources or expertise; 77% report excessive expectations. 1,600 CISOs across 16 countries, fielded May 2026 by Censuswide. Vendor-commissioned.
- The Conference Board (Jul 28, 2026): 55.1% of workers use generative AI or AI agents daily or weekly; 33.3% used employer-provided AI training in the past six months; 28.3% report their organization provides no AI training. Approximately 1,300 workers globally.
- BCG, "AI at Work 2026: Strategy Matters More Than Tools" (Jun 3, 2026): a clear AI strategy lifts measurable business impact by 25 percentage points versus 5 points for better tools alone; 47% spend more time managing AI than doing their actual work; 67% of regular users report improved job satisfaction. 11,749 workers across 14 markets, fielded June 2026.
- Grant Thornton, "2026 AI Impact Survey": 48% of boards have not set AI governance expectations and 46% have not integrated AI risk into ongoing oversight, while 75% have approved major AI investments; 78% lack confidence they could pass an independent AI governance audit within 90 days; COO concern about agentic AI compliance 54% versus 20% for CIOs and CTOs; workforce readiness 29% (CIO/CTO) versus 6% (COO). 950 business leaders, fielded Feb 23–Mar 18, 2026.
- Logicalis, "2026 CIO Report" (Mar 3, 2026): 62% of CIOs report compromising on governance due to limited knowledge; 44% say they fully grasp the risks of AI adoption; 89% describe their approach as "learning as we go." 1,000+ CIOs, fielded by Vanson Bourne. Vendor-commissioned.
- Okta, "Global CISO Insights 2026" (Jul 29, 2026): only 31% of CISOs feel fully aligned with their C-suite and board on acceptable AI risk levels. 306 CISOs and senior security executives; note only 14% of respondents were US-based. Vendor-commissioned.
- ISACA, "2026 AI Pulse Poll" (May 5, 2026): 38% of organizations have a formal comprehensive AI policy; 38% are confident in their board's understanding of AI risk; 22% report AI ROI has met or exceeded expectations; 33% train all employees on AI. 3,400+ digital trust professionals.
- Deloitte Global Boardroom Program, "Governance of AI" (2nd edition): 66% say their boards have limited to no knowledge or experience with AI. 695 respondents across 56 countries, fielded Jan–Feb 2025.
- Data & AI Leadership Exchange, "2026 AI & Data Leadership Executive Benchmark Survey" (Randy Bean, 15th annual): 93.2% identify culture and change management as the greatest impediment; 6.8% cite technology. Approximately 110 senior executives, invitation-only and self-selected — directional only.
- ISC2, "2025 Cybersecurity Workforce Study" (Dec 2025): 48% exhausted from trying to stay current; 47% overwhelmed by workload expectations; 34% report appropriate staffing levels; AI is the top skills gap at 41%. Counterweight from the

same study: 68% report overall job satisfaction and 78% plan to stay in the field. 16,029 respondents, fielded Jul–Aug 2025.

- Splunk and Cisco, "The CISO Report" 2026 (Feb 24, 2026): nearly two-thirds of security teams experience moderate to significant burnout. 650 CISOs across 9 countries, researched by Oxford Economics. 2025 edition (Jan 23, 2025): 29% of CISOs say they receive proper budgets versus 41% of board members who believe budgets are adequate. 600 respondents including 100 board members.
- Dice, "2026 Tech Sentiment Report" (Feb 10, 2026): burnout rose from 31% to 46% of respondents. 1,159 US tech professionals, fielded Nov–Dec 2025. Dice is a technology job board.
- Stack Overflow, "2025 Developer Survey" (Dec 29, 2025): trust in AI accuracy fell to 29% from 40%; 66% report spending more time fixing "almost-right" AI-generated code. 49,000+ respondents.
- data.world and DataKitchen, "2021 Data Engineering Survey" (Oct 19, 2021, Wakefield Research, 600 data engineers): 91% report frequently receiving requests with unrealistic or unreasonable expectations. Cited here only as historical context — it is five years old, predates generative AI, and was commissioned by two vendors selling DataOps tooling.