cognizant®

New frontiers, new *friction*

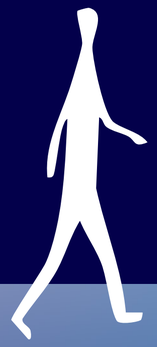

Closing the AI execution gap:
A $2 billion business boost

Table of contents

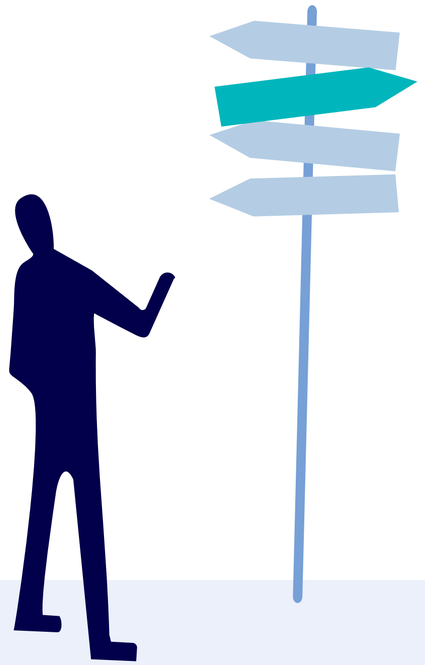

cognizant

In our recent research, two-thirds of leaders have yet to demonstrate measurable business productivity gains from AI, and one-quarter have paused or abandoned deployments.

However, those with a mature tech infrastructure and a focused AI investment are separating from the crowd, and the difference between these companies and others is worth billions in annual returns.

Executive summary

The gap between AI expectations and outcomes is a central concern for business leaders. In our most recent research, only 32% of senior executives can point to tangible gains in business productivity, and 41% don't expect to move the needle on productivity for years. Even when it comes to worker productivity, the numbers are only slightly better, with just over half reporting noticeable improvement.

Meanwhile, one-quarter of organizations have already paused or discontinued AI deployments—**initiatives that were funded, staffed and underway.** We estimate the average company is leaving $4 million[1] in discrete AI investment on the table due to these scrapped projects, this year alone.

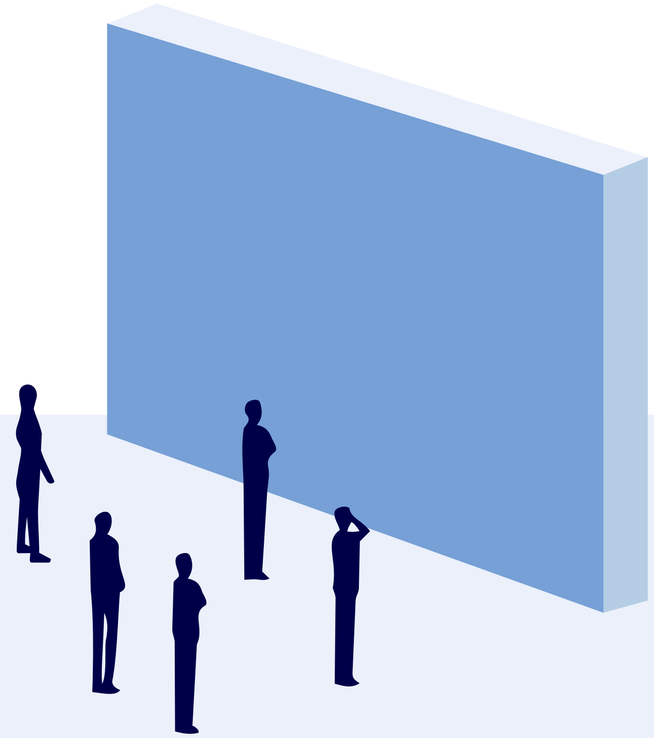

[1] Calculated by applying a 25% attribution rate to the average annual AI investment reported by respondents who had paused or discontinued deployments. This represents an estimate of the portion of investment rendered unproductive by deployment failure and should be treated as indicative.

cognizant

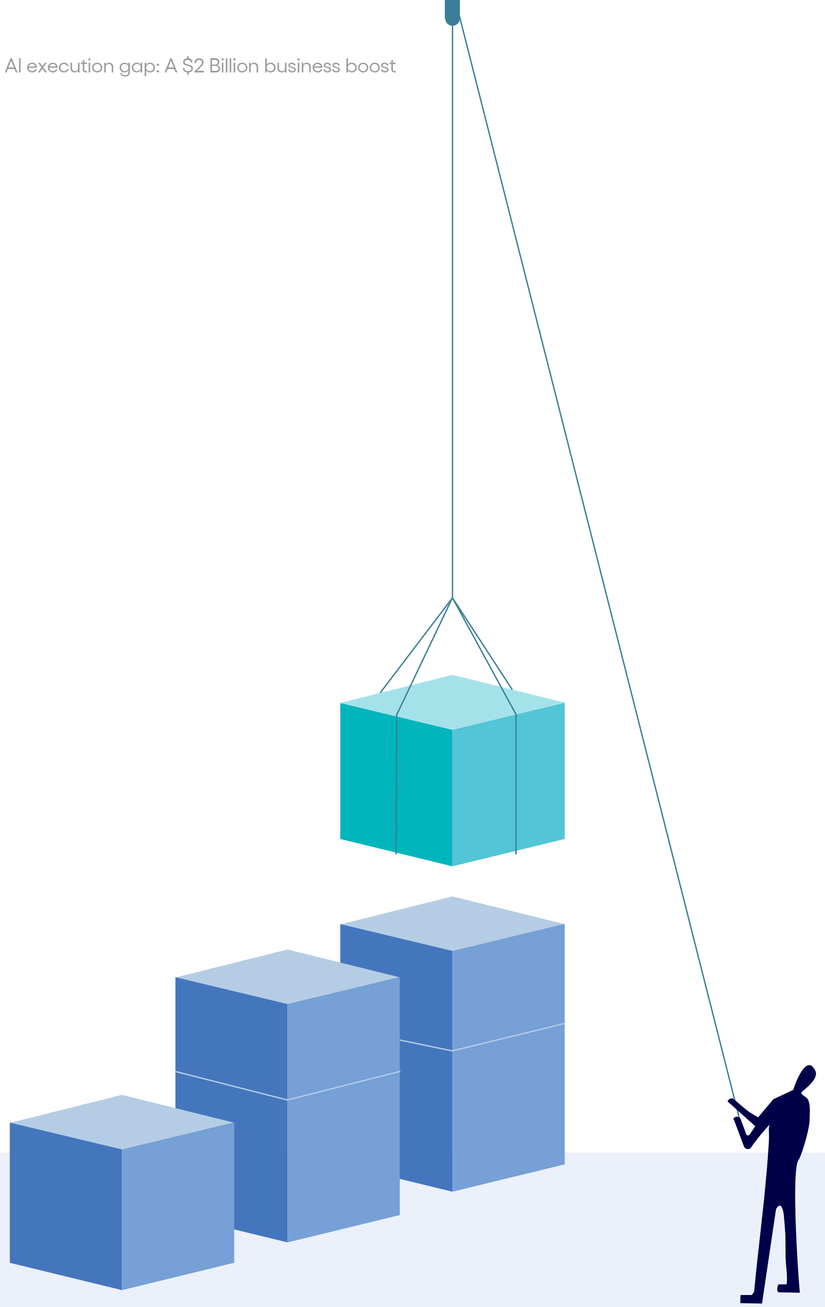

With all leaders in our study expecting to ramp up their AI investment in the next two years, closing this gap matters more than ever. Just under half admit they are not effectively linking their AI spending to specific business value or ROI.

However, some organizations are doing better than others when it comes to realizing AI value. In stark contrast to the plug-and-play AI narratives abundant today, our research reveals that organizations with a mature technology foundation and focused investment on AI technology fundamentals anticipate, and in some cases are already achieving, greater results than those with less developed foundations and a more broad-based AI investment approach.

These outcomes are visible across the most important business measures: business and worker productivity, revenue gains and cost savings.

Our key findings include:

The difference between the lowest and highest performers is dramatic. If a typical G2000 company moved from the weakest to the strongest performing segment, it would realize an estimated $1 billion to $2 billion in annual returns.

Fragile infrastructures can't withstand misaligned AI investment. Organizations that favor non-tech-related AI investments (such as product innovation or talent acquisition) before first bolstering their tech infrastructure are 60% more likely to discontinue an AI deployment than those with similarly weak infrastructures but that choose to invest in AI tech fundamentals first (such as computing power and data readiness).

Even one weak link in the infrastructure chain puts AI outcomes at risk. Businesses that rated all their tech infrastructure dimensions as "good" or "excellent" are far more likely to anticipate and achieve higher gains in all the outcome measures recorded. The gains quickly drop for those with even one dimension rated as "adequate" and fall lower when at least one dimension "needs improvement."

External partnerships are a habit for high performers. Companies focused on investing in AI fundamentals first are the most likely to be working with external partners to do so. In fact, the segment of businesses seeing the strongest results are also those that are most engaged with external partners.

This analysis is part of our series of studies exploring the numerous elements necessary for closing the gap between AI's technical capabilities and real-world results. Our research is based on a survey of 1,100 senior business leaders at G2000 companies and 100 startups across 10 industries.

In this report, leaders will learn how they can achieve significantly better outcomes with AI when they build on a mature tech infrastructure and focus their AI investments, starting with compute, data management and customized AI platforms.

By understanding where to point their AI investments—and the underlying work involved with realizing AI value—**business leaders can move the needle on ensuring outcomes with AI.**

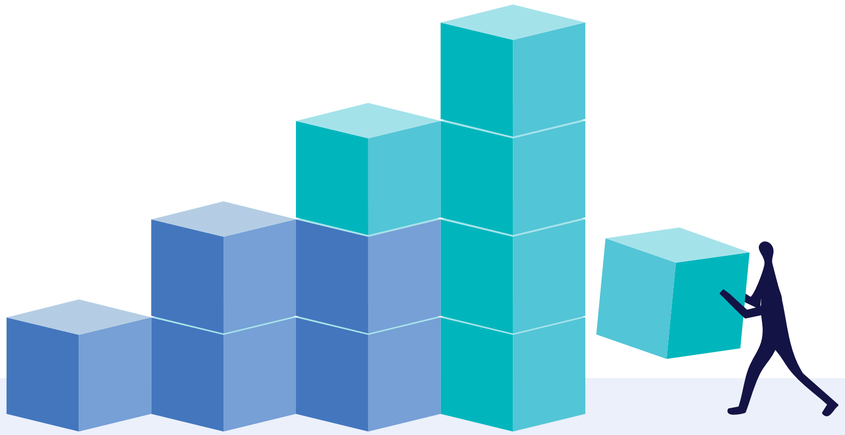

Tech maturity and focused investment drive AI outcomes

To better understand the role of tech infrastructure in realizing gains from AI, we divided respondents into four groups based on the maturity of their tech infrastructure and whether their AI investment was “focused” or “broad-based.”[2] To determine their infrastructure maturity, we scored respondents according to how they rated 10 dimensions of their technology foundation, from “nonexistent” to excellent.[3]

Organizations are classified as “mature” when their average score exceeds the sample median. Those scoring at or below the median are classified as “immature.”

2 Respondents were asked to identify their largest AI investment area. The options we categorized as “focused” were technology infrastructure and customizing or building platforms and tools that enable AI adoption. The options we categorized as “broad” include risk, compliance or governance frameworks; AI product development or innovation; workforce enablement; external industry partnerships; and talent acquisition.

3 Respondents were asked to rate their organization’s current capabilities across 10 infrastructure dimensions: data accessibility, data security, data quality and cleanliness, data traceability and lineage, data reliability and trust, compliance with company rules, compliance with government regulations, compliance with customer privacy requirements, on-premises compute and cloud compute.

cognizant

Higher maturity, focused investment lead to stronger AI outcomes

Infrastructure: Mature ↑ / Immature ↓

Mature

Mature + Broad
233 of 1,100 respondents

Composite score	**12.5%**
Worker productivity	15.3%
Business productivity	12.7%
Cost reduction	10.9%
Revenue increase	10.9%
Deployments paused	22.7%

Best outcomes

Mature + Focused
261 of 1,100 respondents

Composite score	**13.2%**
Worker productivity	17.0%
Business productivity	14.2%
Cost reduction	11.3%
Revenue increase	10.2%
Deployments paused	26.8%

Weakest outcomes

Immature + Broad
299 of 1,100 respondents

Composite score	**9.7%**
Worker productivity	11.2%
Business productivity	11.2%
Cost reduction	8.4%
Revenue increase	8.0%
Deployments paused	32.1%

Building foundations

Immature + Focused
307 of 1,100 respondents

Composite score	**11.4%**
Worker productivity	13.8%
Business productivity	13.4%
Cost reduction	9.3%
Revenue increase	8.9%
Deployments paused	19.9%

Broad ← **Investment** → Focused

We then created an "outcome score"—a composite average of the four outcomes measured: worker productivity, business productivity, revenue gains and cost reduction. The score is based on respondents' self-reported assessment of the outcomes they anticipate or have achieved with AI.

The results are clear: The best outcomes come from organizations with high infrastructure maturity and a fundamentals-first (focused) investment strategy (see Figure 1).

Figure 1
Base: 1,100 senior business leaders at G2000 companies and 100 startups across 10 industries.
Source: Cognizant

Note: Organizations are classified as "mature" when their average score across 10 infrastructure dimensions exceeds the sample median. Those scoring at or below the median are classified as "immature." Investment focus is classified as "focused" when the respondent's primary AI investment area is in technology infrastructure or customizing and building AI platforms. All other investment areas are classified as "broad." The outcome data is calculated from responses to results seen across cost reduction, revenue increase, worker productivity and business productivity. All measures are drawn purely from self-reported survey responses, so they reflect perception of improvement rather than independently verified figures.

cognizant

In fact, organizations in the highest quartile achieve an outcome score that is 3.5 points higher than the lowest quartile, a 31% difference. Further, the highest-performing cohort beats the weakest in all outcome measures, with an extra 5.8 points in worker productivity, 3 points in business productivity, 2.2 points in additional revenue generation and 2.9 points in decreased costs. Crucially, 70% of those in the highest-performing segment are already seeing some of these results, compared with just 42% in the weakest segment.

While these gaps might seem small, when applied to the G2000's combined revenue and cost base of nearly $100 trillion, it represents an estimated $2.5 trillion in unrealized annual value on revenue and cost alone. That rises to $4.7 trillion when worker and business productivity are included. So, for a typical G2000 company, moving from the **weakest to the strongest performing segment is worth an estimated $1 billion to $2 billion** in annual returns.[4]

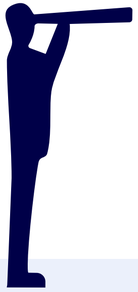

The best outcomes come from organizations with high infrastructure maturity and a fundamentals-first (focused) investment strategy.

4 All figures are indicative of scale rather than precise forecasts. The analysis is based on respondents' self-reported perceived gains on four forms of ROI: worker productivity, business productivity, revenue generation and cost reduction. The dollar figures are derived by applying the gap between our best- and worst-performing segments to the G2000's combined financial base—approximately $52 trillion in revenues and $48 trillion in costs. Labor costs are estimated at 25% of revenues, which is a reasonably conservative assumption across the G2000's mix of capital-intensive and service-oriented sectors. Worker and business productivity carry some overlap with revenue and cost outcomes, which is why we present both a two-dimension figure ($2.5 trillion) and a four-dimension figure ($4.7 trillion). The true value likely sits between the two.

cognizant

Highlights of our findings include:

Mature infrastructures raise the bar across the board.
Tech infrastructure maturity is strongly associated with the breadth and consistency of AI returns, regardless of whether the AI investment is focused or broad. Organizations with a mature infrastructure achieve an average worker/business productivity score of 12.8%, compared with 10.5% for those with immature foundations—a 21% gap. They are also 20% more likely to already be seeing measurable impact from their AI programs. Strong foundations raise the floor for every investment strategy.

However, maturity without focused investment leaves outcomes on the table.
The mature/broad investment group is experiencing a lower outcome score (12.5% vs. 13.2%) than the mature/focused segment. That suggests maturity on its own is not enough: Without directing spend toward the bottlenecks that are most apt to constrain AI at scale, even well-prepared organizations dilute the return on otherwise strong foundations.

Organizations with a mature infrastructure are 20% more likely to realize measurable impact from AI.

cognizant

Consider that most organizations in our study are intent on moving beyond standard AI-powered productivity tools (53%), toward the more scaled endeavors of intelligent automation and operations (62%) and AI-driven analytics and decision support (66%). These are the AI initiatives necessary for moving from worker productivity to the higher-impact measure of business productivity, which requires deeper systemic change.

But to get there, the fundamentals of infrastructure and platforms need to be fortified with investment. Indeed, the mature/focused group leads the mature/broad segment on both worker productivity (17.0% vs. 15.3%) and business productivity (14.2% vs. 12.7%).

In fact, a focused investment strategy pays off at any maturity level.

The bottom half of the matrix reveals another key distinction. While there's a bigger gap in outcomes between organizations with immature versus mature tech infrastructures, the existence of a focused AI investment strategy continues to deliver a boost. One reason for this is that the focus on investing in AI fundamentals can help overcome infrastructure gaps.

Consider that organizations in the low-maturity/focused quadrant achieve an outcome score of 11.4%—much higher than the low maturity/broad group's 9.7%.

Mature/focused group leads the mature/broad segment on both worker productivity (17.0% vs. 15.3%) and business productivity (14.2% vs. 12.7%).

cognizant

Meanwhile, the low-maturity/broad segment has the lowest individual outcomes in every measure, several percentage points below its low-maturity peers, trailing worker productivity by 2.6 points alone. These organizations are investing in the wrong things before the ground beneath them is stable enough to support it.

Top performers are most likely to seek outside help.
Another critical finding is the actions respondents are taking to ensure they close the gap between AI implementation and AI value. The companies with a focused investment strategy are also the most likely to be working with external partners to do so.

In fact, low-maturity/focused organizations are the most likely to work with an external partner, with 75.6% saying this is the case, followed closely by the mature/focused group, at 72.4%. Both groups are well ahead of the two segments with broad investment strategies, at 59.7% for the mature group and 54.2% for the low-maturity group.

The pattern is consistent: The segment seeing the strongest results is the most engaged with external partners, and even those still early in building their foundations are reaching for external help at the highest rate. This suggests that partnership is not just a marker of success but also a key part of how organizations get there.

Organizations with the strongest results are more than 18 points more likely to work with an external partner.

Fundamentals-first investment strategies are a safeguard against failed deployments.
Investment focus becomes the critical variable for organizations still building tech infrastructures. Among the two groups with immature infrastructure, those with focused AI investment strategies have a deployment discontinuation rate of 19.9%.

Those with broad, nontechnical investment priorities sit at 32.1%—making them 60% more likely to pause or abandon a deployment, even before their infrastructure is fully mature. The same group is also 37% less likely to already be seeing results from their AI programs (42.1% vs. 67.1%).

The low-maturity/focused group's discontinuation rate is also the second lowest of all the quadrants, topped only by the mature/broad group. One explanation for the relatively higher discontinuation rate of the mature/focused group is that—given their greater fundamentals-first mentality—they are more likely to take a fail-fast approach to AI experimentation.

Immature/broad organizations are 60% more likely to pause or abandon a deployment than the immature/ focused group.

20%

Organizations with a mature infrastructure are 20% more likely to realize measurable impact from AI.

72% vs 54%

The segment with the strongest results is more than 18 points more likely to work with an external partner.

60%

Immature/broad organizations are 60% more likely to pause or abandon a deployment than the immature/ focused group.

Where infrastructure is falling short—and where to start

Most organizations in our survey do not believe their infrastructure is weak or needs improvement—they rateit as adequate. While this may have been good enough historically, from an AI performance perspective, it can be a showstopper to realizing returns.

cognizant

Consider that organizations with all 10 infrastructure dimensions rated as good or excellent achieve an average gain across both business and worker productivity of 15.6%. When just one dimension is rated as adequate, that drops to 14.1%. This falls further to 12.5%, when at least one dimension is rated as needing improvement.

For previous waves of enterprise technology, adequate systems could often be worked around. Slow access, fragmented data, manual reconciliation or limited compute might reduce efficiency, but they did not necessarily prevent the system from functioning. **AI is less forgiving.** Its value depends on repeated access to large volumes of high-quality data, fast processing, clear provenance, secure integration and continuous use inside live workflows.

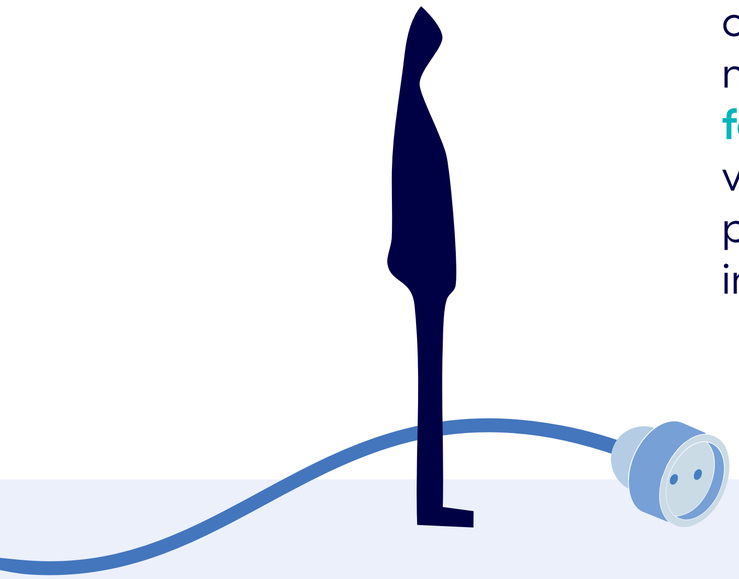

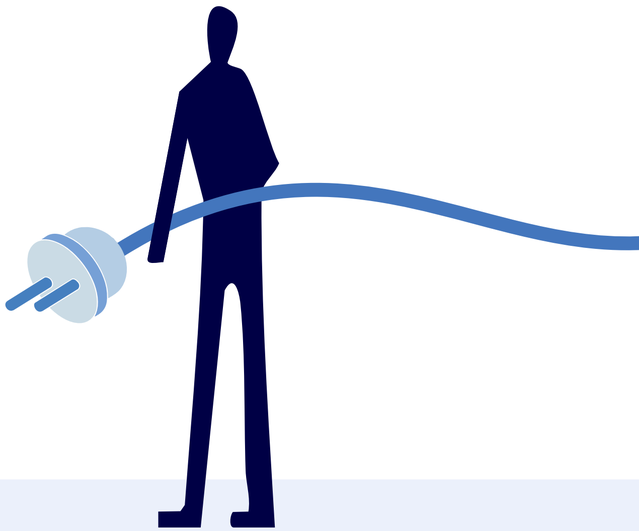

The issue is that many organizations' tech foundations were built for analytics, automation and reporting, not for AI systems that generate outputs, trigger actions and sit inside critical workflows. **Simply put, adequate will no longer suffice, and weak foundations are a compounding drag across the entire AI program.** That said, not all infrastructure challenges are equal in their addressability. A useful way to prioritize is to look at how organizations are currently rating their infrastructure elements and the productivity uplift they can realize from improving on it.

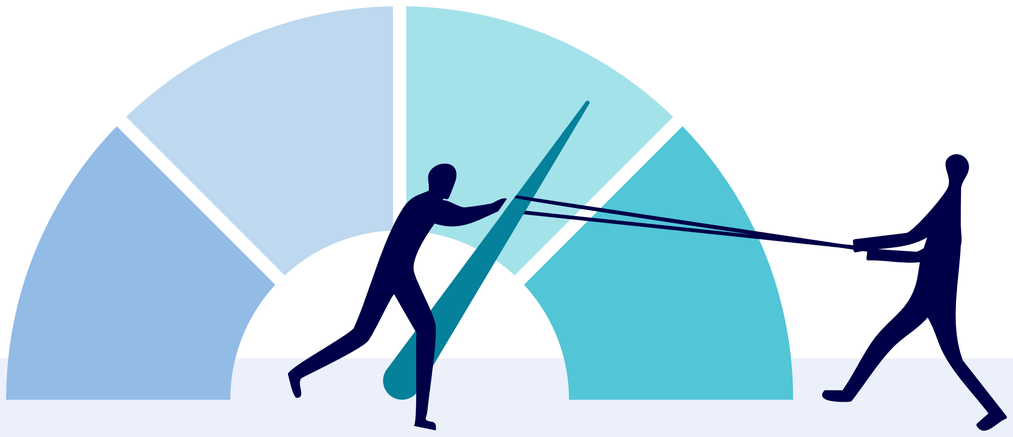

Organizations with all 10 infrastructure dimensions rated as good or excellent achieve an average productivity gain of 15.6%. When just one dimension is adequate, that drops 1.5 points. It falls another 1.6 points when at least one dimension is rated as needing improvement.

On this basis, two infrastructure dimensions stand out:

1) Compute availability: The weakest infrastructure element

One in three leaders cite computing power as an inhibitor for AI adoption. While compute does not create productivity on its own, it does determine whether AI can move from occasional assistance to a normal part of how work gets done.

A team may have access to powerful models, but if employees are restricted in how much they can use them, how much context they can provide, how many documents they can process or how easily AI can be embedded into everyday systems, the productivity benefit suffers.

This bears out in our findings. Of all 10 infrastructure dimensions measured, on-premises compute is the weakest by the widest margin. It is the most likely of all adequate, and only 19.9% rated it as excellent.

The productivity consequences are steep. Organizations with excellent on-premises compute achieve a worker productivity gain that is 3.1 points greater than those with an adequate ranking (16.8% vs. 13.7%).

The gap is even wider with cloud compute. Organizations that consider their cloud compute as excellent outperform those with adequate rankings by 4.8 points (16.7% vs. 11.9%) in terms of worker productivity gains.

The important point is not whether one or other approach (on-premises vs. cloud) is inherently preferable.

In both cases, the underlying issue is the same: Inadequate compute limits AI throughput. Further, compute maturity gives organizations more choice. They can decide when to use frontier cloud models and smaller specialist models, when to deploy open-source models privately and when to combine these approaches. Productivity does not always come from the most powerful model but from being able to match the right model to the right workload at the right cost, latency, context window and level of control.

2) Weak data foundations undermine AI returns

Across the five data-related infrastructure dimensions in our analysis—reliability, security, traceability, accessibility, quality—each creates a different kind of drag on performance. What they share is that 64.5% of organizations have at least one of them rated at only adequate or below. That means most G2000 organizations are running AI programs on data infrastructures that were not built with AI in mind. However, organizations with a strong data foundation (across all data dimensions) report nearly 27% higher productivity gains from AI and are over 20% less likely to discontinue AI initiatives compared with those still working to improve these dimensions.

Organizations with a strong data foundation report nearly 27% higher productivity gains from AI and are over 20% less likely to discontinue AI initiatives.

Here is a look at what businesses can gain by improving their data strategies, in order of where the biggest productivity uplift can be achieved:

Data reliability: When AI systems cannot reliably access the data they need due to unstable data pipelines, model outputs degrade. This is frequently misread as a model problem or a tooling problem, but in practice is usually a data architecture issue.

Among businesses rating their data reliability as needing improvement, the deployment abandonment rate reaches 38.7%, nearly double the rate seen when reliability is rated as good.

Data security and privacy: Nearly one in two leaders identified data privacy and security as a key constraint. The gains that can be realized by improving this infrastructure element are high.

Organizations that rate their data security as good achieve 40% higher productivity gains and are nearly 40% less likely to discontinue AI initiatives than those that rate it as needs improvement.

Data traceability: Of organizations with excellent data traceability, supported by continuous audit trails and strong meta data practices, 45% achieve more than 20% worker productivity gains.

Among those where traceability needs improvement, that figure is 10%. When teams struggle to trace model decisions or stitch together the lineage for an AI model's input data, trust erodes rapidly, and productivity can dive by as much as 33%.

Data accessibility: One-third (33%) of organizations rate their data accessibility as merely adequate. Productivity suffers as a result. The difference in productivity between those with excellent and adequate data accessibility rankings is a full 3 points (15% vs. 12%). Further, organizations with adequate data accessibility are 9 points more likely to have paused or abandoned AI deployments than those with good accessibility (29% vs. 20%).

Data quality: Data quality and cleanliness round out the list, with 25% of organizations at adequate and a 1.8-point productivity uplift available if they improved to a rating of "good." **Data quality at AI scale requires active curation, schema discipline and feedback loops** from model performance back into data pipelines, all of which are essential when creating a remediation path.

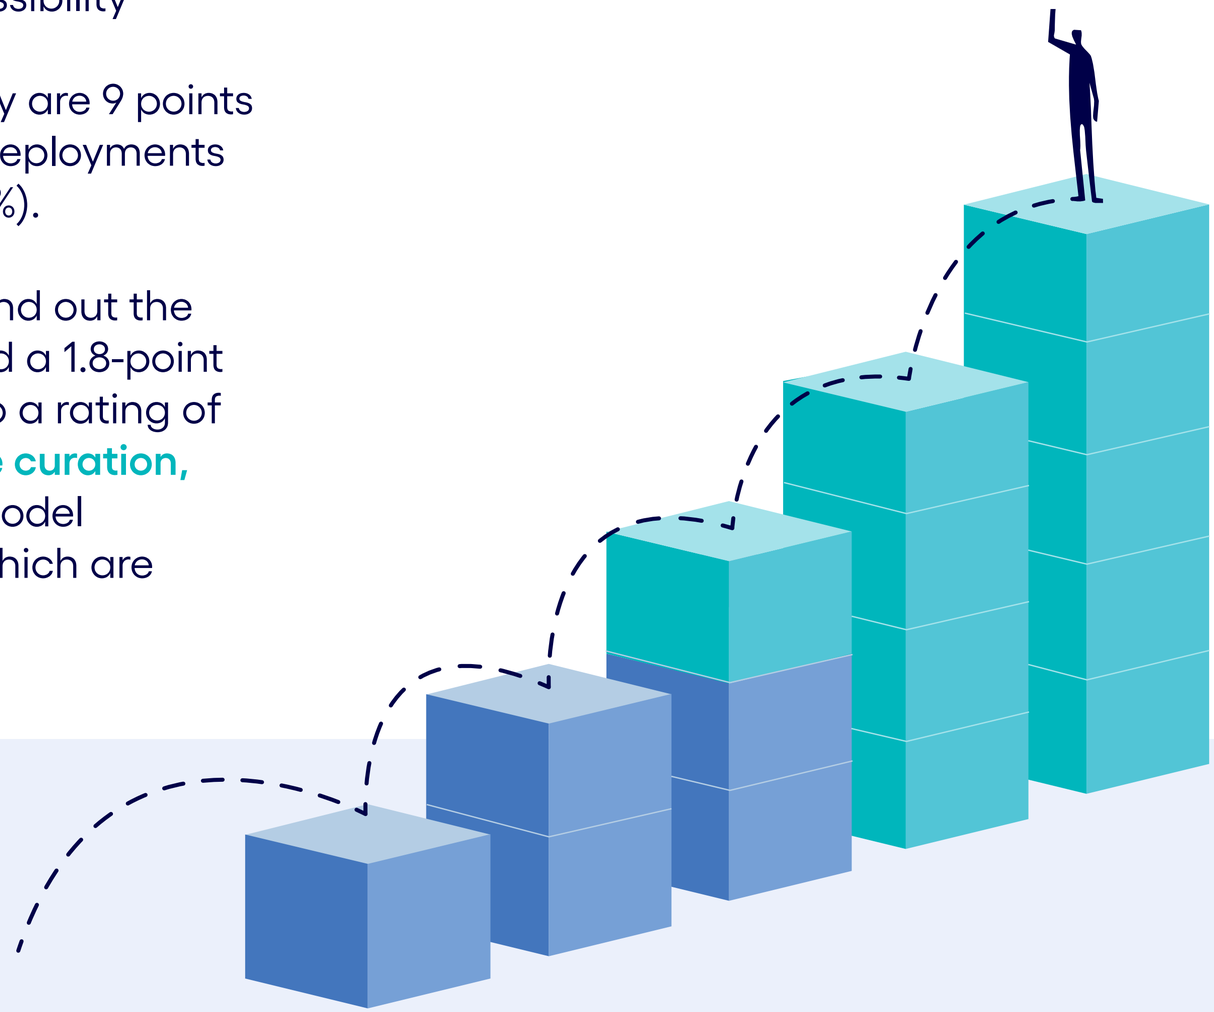

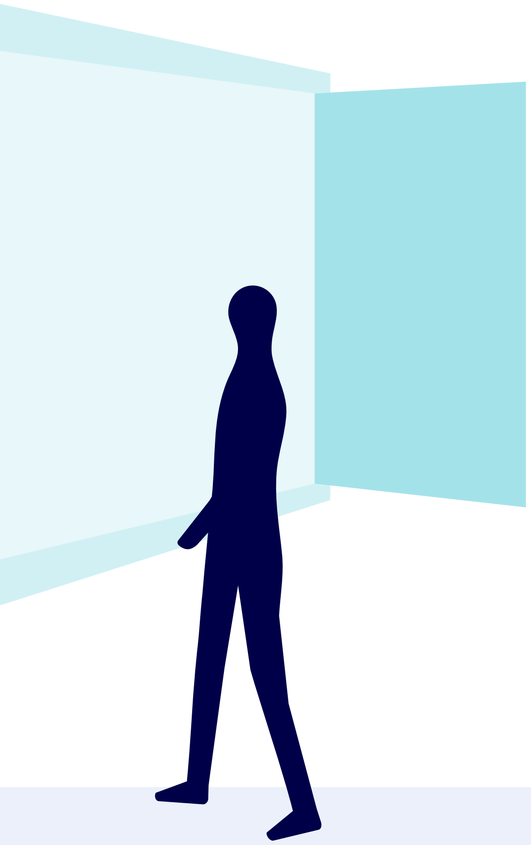

Closing the AI gap: Foundations first, and returns follow

The data is consistent across every dimension we measured: Foundations determine returns. The organizations seeing the strongest results are not those spending the most or moving the fastest. They are the ones treating infrastructure as a prerequisite—investing in compute and their data foundations—before scaling AI capabilities and reaching for external expertise to close the gaps they cannot close alone.

For most G2000 organizations, **the window to course-correct is open but narrowing**. Two-thirds have yet to demonstrate measurable business productivity gains, and a quarter have paused or abandoned deployments. Continued misalignment represents billions in unrealized value and a widening gap between those who build correctly and those who build quickly.

In the end, the infrastructure-building decisions made today will define the AI performance outcomes of tomorrow.

Authors

Ollie O'Donoghue
Head of Cognizant Research

Rijo Thomas
Director, Cognizant Research

Duncan Roberts
Associate Director, Cognizant Research

Alexandria Quintana
Senior Manager, Cognizant Research

William Brandel
Research Advisor

cognizant®

Cognizant (Nasdaq: CTSH) is an AI Builder and technology services provider, bridging the gap between AI investment and enterprise value by building full-stack AI solutions for our clients. Our deep industry, process and engineering expertise enables us to build an organization's unique context into technology systems that amplify human potential, drive tangible outcomes and keep global enterprises ahead in a fast-changing world. See how at cognizant.ai or @cognizant.

World Headquarters

300 Frank W. Burr Blvd.
Suite 36, 6th Floor
Teaneck, NJ 07666 USA
Tel: +1 201 801 0233

European Headquarters

280 Bishopsgate
London
EC2M 4AG
England
Tel: +44 (01) 020 7297 7600

India Corporate Office

Siruseri-Software Technology Park of India (STPI)
SDB Block—Ground Floor North Wing
Plot No H4, SIPCOT IT Park
Chengalpattu District
Chennai 603103, Tamil Nadu
Tel: 1800 208 6999

APAC Headquarters

1 Fusionopolis Link, Level 5
NEXUS@One-North, North Tower
Singapore 138542
Tel: +65 6812 4000

WF - 4811000