


WHITEPAPER

The Adoption Advantage

Why Trusted Data, Not the Next Model, Will Determine Enterprise AI Returns




Content

Executive summary	2
What the evidence shows	3
1. The pattern: performance races ahead, then becomes good enough	4
2. Physical ceilings: when heat, chemistry or regulation sets the limit.....	5
2.1 Smartphone cameras: megapixels surged, then plateaued	5
2.2 Phone batteries: capacity climbed, but all-day life was already enough.....	6
2.3 Processors, storage and memory: the same plateau pattern	7
3. Software reaches human ceilings, not physical ones	8
3.1 Speech recognition reached human parity, then competition shifted	8
3.2 Machine translation, audio and features: the same ceiling, everywhere	9
4. Frontier AI is crossing the same threshold	10
4.1 Usage is surging: the market has already decided	11
4.2 Data readiness is the entry ticket, not a later phase.....	12
4.3 Data governance and AI governance are different jobs	13
4.4 AI has hardware ceilings too: chips, data centres and power.....	13
4.5 The industry itself is now arguing for a slower frontier.....	13
5. Winning the adoption race.....	14
5.1 Run a portfolio, not a single bet.....	14
6. The cost of waiting rises every month	15
6.1 When waiting is still the right call.....	16
7. What leaders should do now.....	17
Conclusion.....	18

Executive summary

Every headline about artificial intelligence is about a bigger number, a smarter model, a higher score, a new frontier. That race is real and it is still running albeit slowing. But for a large and growing share of everyday business work, the next increment of frontier intelligence is not what stands between a company and results. What stands in the way, is adoption. The unglamorous work of connecting AI to trusted data, wiring it into workflows, governing it, and teaching people to use it.

Here we make the case through a pattern every leader already knows. Technologies race on one specification and then cross a “fit for purpose” line, after which the specification sheet stops selling the product and value shifts to how the technology is used. There are two versions of that ceiling. Hardware like phone cameras, batteries, processors, disks, etc. stops at a physical wall such as heat, chemistry, and even regulations. Software like speech recognition, translation, audio compression, etc. keeps improving but its value flattens at a human ceiling, i.e. the point where more capability exceeds what a person can perceive or a task needs. Frontier AI is software at that stage. For a huge range of business tasks, it is now crossing its human ceiling, exactly as speech recognition and translation did before it.



AI is still advancing, but for most business tasks it has already cleared the bar. The contest is no longer who has the smartest model, but who deploys capable models best. That advantage is built, not bought and every month on the sidelines is a month not building it.

What the evidence shows



Hardware hit physical walls. Phone cameras sat at 12 megapixels for seven years; CPU clock speed has been near-flat since 2004; batteries stalled at a ~5,000 mAh ceiling with real endurance barely moving for a decade.



Software hit human ceilings. Speech recognition reached human-transcriber parity (~5.1% error) in 2016-17 and flattened. Machine translation “bridged the gap” and became good enough for everyday use. In both, the race moved from accuracy to deployment.



Demand has already voted. Google alone went from 9.7 trillion tokens a month in 2024 to over 3.2 quadrillion in 2026. Businesses do not spend on technology at such scale that is not yet useful.



The bottleneck has moved. Dun & Bradstreet’s survey of 10,000 businesses in June 2026, finds 97% of businesses with active AI initiatives globally, yet only ~4% of Indian businesses say their data is fully ready for AI at scale.



Capability is not the only constraint. Power, grid connection and cooling; not just chip supply are now the main brake on AI infrastructure, and the leading labs are publicly arguing to pace the frontier. Both point the same way - build on the capability, enough for the task at hand.



Waiting compounds against you. Value from a general-purpose technology is unlocked through organisational learning, not at purchase. A better model handed to a latecomer does not come with the data, workflows and skills the early mover has already built.

Waiting for the next model before committing is like refusing to pull out of the driveway until every traffic light in the city is green. The lights never go all green at once and while you wait, everyone else is already

moving. AI technology is already deployment-ready to create value today and advantage now goes to organisations that adopt, fix their data, redesign their workflows, learn and deploy faster than competitors.

1


The pattern: performance races ahead, then becomes good enough

Mass technology tells the same two-act story. In the first act, a single headline number climbs fast and sells products like megapixels, gigahertz, milliamp-hours, reducing error rates. In the second act, that number either stops climbing or keeps climbing while the value to a normal user flattens. The commercial turning point is not when

technology stops improving. It is when it becomes good enough, because that is when the buyer stops caring about the specs and starts caring about the experience.

Crucially, there are two kinds of ceiling, and telling them apart is what makes this analogy hold for AI rather than break on it.

Two kinds of ceiling - same business result



Two ceilings, one shift. Hardware hits physical limits: software keeps improving until extra capability adds little practical value. In both cases, value moves from better specifications to better adoption.



2

Physical ceilings: when heat, chemistry or regulation sets the limit


The most familiar plateaus are hardware, because physics eventually says no. Three short examples make the point.

2.1 Smartphone cameras: megapixels surged, then plateaued

The iPhone camera is the cleanest example. Resolution raced from two megapixels to twelve megapixels in eight years, then sat at twelve megapixels for seven years (2015-2022) while Apple competed on stabilisation,

computational photography and extra lenses instead. It stepped once to forty-eight megapixels in 2022 and has stayed there since. Industry-wide, average resolution rose from twenty-seven megapixels in 2020 to 54 in 2024, yet even 200-megapixel sensors are routinely binned down into cleaner 12 or 50-megapixel images. Beyond a threshold, the bigger sensor number is still sold, but the value is routed into image quality and low-light performance rather than raw resolution.

Smartphone camera: iPhone megapixel race, then plateau




Two different ceilings, one outcome: once capability passes what the task or the human can use, more capability stops creating value.

Sources: Apple model specifications; Counterpoint Research (54MP average, 2024)

2.2 Phone batteries: capacity climbed, but all-day life was already enough

Batteries make the same point from a different angle. Flagship cell capacity climbed from 1,400 mAh in the first iPhone (2007) to about 5,000 mAh today, but lithium-ion energy

density has improved only around 5% a year over the past decade, and makers cannot simply keep enlarging the cell: a bigger battery means a thicker, heavier, costlier phone and tighter safety and shipping constraints. The result is a soft ceiling near 5,000 mAh for mainstream flagships.



Capacity climbed ~3.5× toward a ~5,000 mAh soft ceiling, but the outcome customers feel — all-day use — was met years ago and held; extra capacity and efficiency were reinvested in richer experiences (screens, 5G, software, AI), not longer life.

Sources: iPhone teardown capacities (2007–2025); independent endurance testing (Tom’s Guide, GSMArena, 2016–2026); Li-ion energy density ~5%/yr (Ars Technica, 2021); Class 9 lithium shipping (US 49 CFR 173.185 / UN3480).

By the standards that matter to most users, all-day battery life was achieved years ago. Yet every improvement in batteries and power efficiency has been offset by bigger displays,


5G, heavier software, and AI features. The numbers on the spec sheet kept rising, while the day-to-day experience remained largely unchanged.

2.3 Processors, storage and memory: the same plateau pattern

Processors tell the story faster. The “power wall” of 2004 made higher clock speeds thermally unmanageable, so clock speed went flat for two decades, about 4 GHz in 2013 to 6.2 GHz in 2024, against a near-doubling every

two years before and the industry stopped selling gigahertz, pivoting to more cores and efficiency. Storage is the mirror image: the specification never stopped climbing, reaching thirty-six terabyte drives by 2025, but individual need plateaued long ago, and the value migrated to the cloud, to access and to use.

Processors hit a wall: the gigahertz race stalled after 2004



The “power wall” froze clock speed for ~20 years; the industry pivoted to more cores and efficiency, not GHz.

Source: Intel/AMD desktop CPU history; [processorhistory.com](https://www.processorhistory.com/).



3


Software reaches human ceilings, not physical ones

A fair objection is that all of those were hardware limits – heat, chemistry, regulation – and that software faces no such wall. That is right, and it is exactly why software is the better analogy for AI. Software stops at a human ceiling: the point where additional capability exceeds what a person can perceive or what the task requires. Past that point the number can keep climbing, but the race moves to convenience, integration and adoption.

3.1 Speech recognition reached human parity, then competition shifted

Speech recognition is the clearest precedent. For decades, progress was measured by word error rate on the Switchboard benchmark. In 2017 Microsoft reached human parity at 5.1%, against 5.9% for professional transcribers. The contest immediately moved from accuracy to deployment with captions, dictation, meeting transcription, and voice interfaces as economics shifted from performance to adoption.

Speech recognition: software hits its human ceiling



The pursuit of human-parity speech recognition culminated around 2016-17 at roughly 5.1% error rates. Once that milestone was reached, innovation moved up the stack, from understanding words to generating captions, powering dictation and enabling conversational assistants.

Source: Microsoft Research (Switchboard), Word Error Rate on the Switchboard conversational-speech benchmark.

3.2 Machine translation, audio and features: the same ceiling, everywhere

The pattern repeats across software. Google's 2016 neural translation system cut errors by around 60%, and once translation was good enough for everyday use the contest moved to embedding it across products and workflows.

Audio compression is the purest case: above roughly 192-256 kbps, listeners cannot distinguish a compressed track from the original, so encoders stopped chasing bitrate because the human ear is the wall. Office software overshot what most people use years ago, which is precisely why simpler, cheaper tools could win on convenience rather than capability.

Performance oversupply

Clayton Christensen described this pattern: technology progress routinely outstrips what customers can use, after which point the basis of competition flips from raw performance to price, convenience and integration. Sufficiency is not a failure of technology, but the moment value moves from the specification to how the technology is used.

What both ceilings teach us

Technology becomes economically transformative when it is good enough, affordable and easy to deploy, not when its headline metric stops improving. Hardware reaches that point at a physical wall; software reaches it at the limit of human usefulness. For many business tasks, frontier AI is crossing that human ceiling now.



4

Frontier AI is crossing the same threshold

Major technologies tend to follow a similar pattern, whether hardware or software. In the early stages, each new generation changes what users can do with it. Later, improvements keep coming but stop changing the outcome for everyday tasks, and value shifts from the headline specification to how the technology is used. Mobile telephony shows this clearly. 2G made mobile voice and text widely available, 3G brought the internet to the phone, and 4G made video, maps, and digital payments work smoothly on the move. 5G is faster again, yet for most people streaming, browsing and messaging feel much the same, because 4G had already cleared the bar for those tasks. Much of 5G's distinctive value is in specialised uses such as industrial automation, private networks and low-latency applications. 6G will push the frontier further, but no one is waiting for it to send a message or make a payment. The same pattern appears in processors, which stopped competing on clock speed after 2004, and in phone cameras, which stayed at 12 megapixels for about seven years. Speech recognition followed it too when it reached human parity.

Frontier AI is now following the same pattern, crossing the threshold one task at a time. Two things are therefore true together. AI progress is not slowing as models keep improving at

reasoning, coding, multimodal and agent tasks. Yet for a large share of business work, today's models are already "good enough", meaning they are past the point where more capability changes the result. They reliably summarise documents, sort enquiries, draft standard replies, extract data and search approved knowledge. On tasks like these, a better model adds only marginal gains, much as GPS that already places you within a few metres gets you to a meeting no faster with centimetre precision, unless you are landing an aircraft. The frontier keeps moving toward harder problems, while routine work has already crossed the threshold, just as everyday mobile use crossed it with 4G.


That is why waiting no longer makes economic sense. Waiting pays only if the next model will be much better at your work. For routine tasks, that extra benefit is now small, and the cost of getting started has collapsed. What does not get cheaper with time is the work that decides returns on technology investments that is clean, connected data, redesigned workflows, trained people and sound governance. Firms can build these only by doing the work. A firm that waits gets a slightly better model. A firm that starts now gets a learning curve, and it is the learning curve that compounds.

4.1 Usage is surging: the market has already decided

The clearest signal that AI is already good enough for many tasks is its scale of use. Google reported processing 9.7 trillion tokens a month across its products and APIs in 2024,

rising to more than 480 trillion in 2025 and over 3.2 quadrillion by 2026. OpenAI reported that more than 9,000 organisations each processed over ten billion tokens. Businesses do not invest on this scale in technology that has yet to prove useful.

AI usage is exploding - aggregate tokens per month



9.7T (2024) → 480T (2025) → 3,200T = 3.2 quadrillion (2026), across Google products & APIs. Not a model-by-model series.

Source: Google I/O opening keynotes, 2025 and 2026.



4.2 Data readiness is the entry ticket, not a later phase


If model quality were still the binding constraint, businesses would be starved of capability. Instead, they are starved of readiness. In India, the pattern is stark. Dun & Bradstreet's AI Momentum Survey finds that every business surveyed has AI projects underway, more than half have moved beyond pilots into production, operations or agentic workflows, and 73% report some measurable return. Yet only 4% say their enterprise data is fully ready to support AI at scale. Most describe their data as partially ready (58%) or mostly ready (30%). That is enough to start, but not enough to scale. What matters now is not the model, but whether the data can support decisions consistently across the enterprise.

Read that gap correctly. Data readiness is not a later phase of an AI programme; it is the entry ticket. A model can only reason over what it can reach, and what it reaches

must be complete, current, correctly resolved and permissioned. Where records are partial, duplicated or stale, the system still returns an answer, but it simply returns a confident one that cannot be trusted, and the failure arrives as a business error rather than a technical one.

There are two different bars, and leaders should hold them apart. To start with today's AI, an organisation needs enough clean, accessible data for a bounded set of tasks: a defined source of truth, current records, clear permissions and a way to check the output. That bar is reachable this quarter for most firms. To leverage AI fully, agents acting across systems, decisions taken with or without a person in the loop – the bar rises to full data sanctity: every system agreeing on the identity of the same customer, supplier or entity, with lineage, refresh discipline and controls behind it. AI agents will use enterprise data continuously and at a scale humans never did, so every system an agent touches must agree on who the entity is. The first bar starts the clock; the second is the work that compounds.

Adoption is high. Data readiness is not.



What stands in the way of that convergence is not the model itself, but the data. Globally, only around 6% of firms consider their data fully AI-ready.

Source: Dun & Bradstreet AI Momentum Survey of 10,000 businesses, across 32 countries (Q3 2026).

4.3 Data governance and AI governance are different jobs

These two are routinely merged in board papers, and the merge is expensive. Data governance answers questions about the asset: who owns this data, where it came from, how current it is, who may use it, how it is corrected, retained and protected. AI governance answers questions about the system's behaviour: whether a use case is permitted, whether the model was evaluated for that task, who reviews the output, how drift is monitored, where the audit trail sits and who is accountable when it is wrong.

The dependency runs one way. AI governance sits on top of data governance and cannot substitute for it: a well-documented model-risk framework applied to untrusted data simply produces auditable bad decisions. Practically, that argues for two linked but separate mandates; a data owner accountable for quality, lineage and permissions, and an AI owner accountable for use-case approval, evaluation and monitoring; meeting at a shared register of which data feeds which system, for which decision.

4.4 AI has hardware ceilings too: chips, data centres and power

AI may be software, but its future growth is increasingly governed by the laws of physics and infrastructure. The industry's constraint has shifted from acquiring GPUs to securing sufficient power. Electricity availability, grid upgrades, transformer shortages and interconnection delays now determine how quickly new AI capacity can be added. As data centre electricity consumption rises toward twice current levels by 2030 and cooling requirements intensify, the economics of AI scaling are becoming constrained by energy as much as by algorithms.

This matters commercially in two ways. It caps how quickly the frontier can scale, which weakens the case for waiting on a step change that compute and power may not deliver on schedule. And it puts a floor under the cost of frontier inference, which strengthens the case for a model portfolio: route routine, bounded work to efficient models and reserve scarce, expensive frontier capacity for the problems that genuinely need it.

4.5 The industry itself is now arguing for a slower frontier

In September 2026 that argument moved inside the labs. Anthropic's Dario Amodei published an essay, "We Must Pace the Frontier", calling for AI companies to deliberately slow the rate at which model capabilities improve so that safety research, third-party evaluation and oversight can catch up. Sam Altman, Demis Hassabis and Elon Musk publicly backed the direction within days, and OpenAI said it would adopt the same embedded independent-evaluator model. Pacing is explicitly not a halt to training or technical progress.

For a business, this settles the question of what exactly one would be waiting for. If the labs themselves intend to slow the rate of capability jumps, the next large step is further away; while the capability already shipped is not going anywhere. That is the definition of a good moment to build on what exists: a stable base, a longer window before the ground shifts, and no reason to believe that a delayed start will be rewarded with an easier one.

5

Winning the adoption race


Once many tasks are past the threshold, the contest is no longer “who has the smartest model?” but “who deploys capable models best?” That is a race about data, workflows, governance and people and, unlike model access, it cannot be bought off the shelf on launch day.

5.1 Run a portfolio, not a single bet

Just as a business does not put its most expensive employee on every routine task, it should not route every AI request to the most powerful and most expensive model.

Economics rewards a portfolio. Leading AI providers offer model families with large cost differences between efficient and frontier tiers. OpenAI’s GPT-5.6 lineup spans roughly a 20x range in input-token pricing, while Anthropic’s Claude family spans about 10x. Rather than defaulting every workload to the frontier, organisations can route routine, repeatable tasks to lower-cost models and reserve frontier systems for problems where additional capability materially changes the outcome. The objective is not to buy the smartest model for every task, but the most cost-effective model that clears the required quality threshold.

Model Portfolio: a 50-60x price range across tiers



60x
output-price spread
\$0.40 → \$25.00

BUDGET Gemini 2.5 Flash-Lite
\$0.10 in | \$0.40 out

MID Claude Sonnet 5
\$2.00 in | \$10.00 out

FRONTIER Claude Opus 5
\$5.00 in | \$25.00 out

The takeaway is not “buy cheap.”
It is “route deliberately.”

Operating model: Route bounded work to efficient models, escalate only the hard tail to the frontier, and treat data, evaluation, integration and governance as the foundation.

6


The cost of waiting rises every month

Waiting for a better model feels prudent, but history suggests it is often the more expensive strategy. General-purpose technologies do not create value at the time of purchase; they create value as organisations learn how to use them. NBER working paper analysing cloud-computing data from ~100,000 firms found that new adopters improve productivity by 33% in their first year and continue improving for four years, indicating that much of the value comes from organisational learning and experience rather than adoption alone. AI will

follow the same curve. A better model handed to a latecomer does not arrive with trusted data, redesigned workflows, trained users or production governance already in place.

The cost of waiting. One firm adopts now and climbs the learning curve; another waits 18 months for “the next model” and starts from zero. The second firm may have a better model, but the first has the operating advantage.

A better model doesn't hand a latecomer the learning curve



Don't idle in the driveway: the early mover climbs the curve; the latecomer starts from zero with a better model but no data or workflows.

Illustrative; shape based on NBER WP 32938 (~33% first-year cloud-productivity gain, stabilising after ~4 yrs).

The traffic-light problem

Waiting for the perfect AI model is like refusing to leave the driveway until every traffic light in the city turns green. That moment never comes. The frontier will keep moving, and while you idle, competitors are already three intersections ahead, learning the route, fixing the bottlenecks and building confidence with every mile. The practical move is to pull out on the first clear stretch and improve as you go.

The head start compounds. Early movers accumulate clean, verified and connected data; redesign workflows that future models can plug straight into; build teams that extract more value from the same tools; and create the governance and trust needed to move pilots into production. When the next model arrives, they are not starting again. They are upgrading an engine already on the road. The latecomer is still assembling the car.

6.1 When waiting is still the right call

“Good enough” is task-specific, not universal. For a defined set of high-stakes work, current reliability is genuinely the limiting factor, and prudent leaders should wait, pilot narrowly, or keep a human firmly in control:

- Safety-critical or irreversible decisions, where errors cannot be reliably detected, reversed or contained.
- Unsupervised high-value actions in finance, law, medicine, or infrastructure.
- Long-horizon autonomous tasks, especially when they span multiple systems and reliability still fails.
- Work with no effective human review, audit trail or rollback.

The rule is simple: deploy when the whole system, not just the model, clears the business threshold; wait only when model reliability is truly what is missing.



7

What leaders should do now

The strategic shift is to stop treating each model release as a reason to delay, and to start competing on adoption. Six actions follow directly from the evidence.

1

Anchor to what today's models already do well. Build the plan around bounded tasks, and treat future model gains as upside, not a precondition.

2

Fix the data foundation first. Only ~6% of firms call their data fully AI-ready; verified, connected, permissioned data is the single biggest block to unlock ROI.

3

Run a model portfolio. Route each task to the cheapest model that clears its quality bar and reserve the frontier for hard high-value work.

4

Redesign the workflow, not just the tool. Value comes from re-engineering the process around AI and measuring outcomes, not bolting a chatbot onto an unchanged process.

5

Govern for production from day one. Put privacy, auditability, human review and model-risk controls in place so pilots can scale.

6

Start the learning curve now. Deploy early to build the data, workflows and organisational capabilities that compound over time. Those advantages must be earned and will not be handed to late adopters by the next generation of models.



CONCLUSION

The camera in your pocket stopped competing on megapixels years ago. Batteries hit a practical ceiling, and processors stopped winning on gigahertz. Yet these devices became more valuable, because the contest moved from the headline number to how the technology was designed, integrated and used. Speech recognition and translation crossed the same line in software. They stopped not at the limit of physics but at the limit of everyday usefulness.

Frontier AI is now crossing that line for a large share of business work. The models will keep improving, but they are also becoming available to everyone, at falling cost. What a competitor cannot rent is the quality of the data a business runs its AI on. That is where advantage now lies.

It matters more because AI is moving from answering questions to taking actions: approving credit, onboarding suppliers, releasing payments and flagging risk. When AI only advised, a data error produced a wrong answer that someone could catch. Once AI acts, the same error becomes a wrong decision, repeated at scale before anyone notices. The foundation has to meet four tests:

- **Trusted:** Every input comes from a known, reliable source. Otherwise, a fluent answer can rest on a stale or unverifiable record.
- **Verified:** Every business is confirmed to exist and to be who it claims to be. Otherwise, AI can just as efficiently onboard a shell company or pay a fraudulent supplier.
- **Interconnected:** The same entity is recognised across finance, procurement, sales and risk systems, and linked to its owners, parents and subsidiaries. Otherwise, exposure to a single group stays hidden across separate systems.
- **Governed:** Every output can be traced to its source and explained to a regulator, an auditor or a board. Otherwise, no one can stand behind the decision.

Weak data does not make AI fail loudly. It makes AI wrong confidently, at speed and at scale. Strong data does the opposite: every new use case reuses the same foundation, so each one is faster and cheaper to deploy than the last. As models become increasingly interchangeable, trusted data emerges as the more enduring source of competitive advantage.

So, data is not something to sort out after the AI investment. It is the first investment, and the one every model will depend on. Do not wait for the green light. Start with the tasks where today's AI already clears the bar, base them on trusted, verified and connected business data, and govern them from the start. Every month spent learning compounds, and every month spent waiting widens the gap. The next model will not close that gap. Firms with stronger data will get more out of it, and the gap will grow.





ABOUT DUN & BRADSTREET®

Dun & Bradstreet, a leading global provider of business decisioning data and analytics, enables companies around the world to improve their business performance. Dun & Bradstreet's Data Cloud fuels solutions and delivers insights that empower customers to accelerate revenue, lower cost, mitigate risk and transform their businesses. Since 1841, companies of every size have relied on Dun & Bradstreet to help them manage risk and reveal opportunity. For more information on Dun & Bradstreet, please visit www.dnb.com.

Dun & Bradstreet Information Services India Private Limited is headquartered in Mumbai and provides clients with data-driven products and technology-driven platforms to help them take faster and more accurate decisions across finance, risk, compliance, information technology and marketing. Working towards Government of India's vision of creating an Atmanirbhar Bharat (Self- Reliant India) by supporting the Make in India initiative, Dun & Bradstreet India has a special focus on helping entrepreneurs enhance their visibility, increase their credibility, expand access to global markets, and identify potential customers & suppliers, while managing risk and opportunity.

India is also the home to Dun & Bradstreet Technology & Corporate Services LLP, which is the Global Capabilities Centre (GCC) of Dun & Bradstreet supporting global technology delivery using cutting edge technology. Located at Hyderabad, the GCC has a highly skilled workforce of over 500 employees, and focuses on enhanced productivity, economies of scale, consistent delivery processes and lower operating expenses.

Visit www.dnb.co.in for more information.

Disclaimer:

Dun & Bradstreet makes no representations or warranties, express or implied, with respect to information provided and the results of the use of such information, including but not limited to implied warranty of merchantability and fitness for a particular purpose, warranties of accuracy, completeness, or currency. Neither Dun & Bradstreet nor any of its parents, subsidiaries, affiliates or their respective partners, officers, directors, employees or agents shall be held liable for any damages, whether direct, indirect, incidental, special or consequential, including but not limited to lost revenues or lost profits, arising from or in connection with use or reliance on the information provided in this report. Recipients of these reports are responsible for determining whether the information contained herein is sufficient for use and shall use their own skill and judgement when choosing to rely upon the reports.

© Dun & Bradstreet, Inc. 2026. All rights reserved.