


Entrepreneur
MIDDLE/EAST
SPECIAL REPORT 2026



THE STATE OF 2026



How the Middle East
is Building the Future
of Intelligence



Powered by:



TABLE OF CONTENTS

Editor's Note	03
The State of AI in 2026	05
AI is Reorganizing the Global Economy – And the Shift Has Already Begun	07
The State of Global AI in 2026	09
Money Flows: What AI Investments in 2026 Tell Us	15
10 Graphs That Tell Us About the State of AI Today	19
The King Pins of AI	31
The Man Trying to Enable a Sovereign Online Future Built by AI	45
AI's New Power Corridor: The UAE and Saudi Arabia	49
Key Themes and Innovations Shaping 2026	53
Copilots: From Adoption to Integration	55
Agentic AI: When Software Starts Deciding and Acting	57
AI in Healthcare: From Clinical Validation to Predictive, Personalized Medicine	61
AI in Financial Systems: Expanding Access, Managing Risk	63
AI and Cybersecurity: Escalation Without Equilibrium	65
AI & PropTech: How Reno App is Using AI to Transform Real Estate's Most Unpredictable Layer	67
The Challenges Ahead in the Global AI Race	71
Dubai's AI Positioning	83
Dubai's Next Bet: Agentic AI as the New Growth Engine	85
How UAE Ranks in the Global Industry 5.0 Index	89
Q&A	91
Dubai AI Campus: The Institutional Bet to Make Dubai a Global AI Launchpad	93
Dubai AI Festival 2026: A Convergence of Capital, Governance and Deployment	95
CONCLUSION: The Infrastructure Decade of Intelligence	97

EDITOR'S NOTE



Erika M. Welch

Erika Masako Welch
Chief Content Officer, Lucidity Insights

Twelve months ago, the question was whether artificial intelligence would reorganise the global economy. That question is settled. The one that remains is far more interesting: who is organised to capture it, and how fast?

It is my privilege to introduce The State of AI 2026, the official publication of the Dubai AI Festival. This has been a year of AI breakthroughs arriving at breakneck speed, which has made keeping up a challenge from a publishing standpoint. A report written in one month, may already be outdated in the next.

Let's start with deployment. Nine in ten organisations now use AI regularly, and nearly one in five working-age adults


worldwide use it. But breadth is not depth. Only 44% of organisations say AI is scaling across the enterprise, just 37% see any effect on earnings, and a mere 6% attribute 5% or more of EBIT to it. The gap between AI activity and AI impact is the defining story of 2026, and it is widening by company size: 40% of billion-dollar firms are now scaling AI agents, against 22% of everyone else.

The agents are coming regardless. On SWE-bench Verified (a standard test of whether AI can do real software-engineering work), the best systems went from solving 4% of real-world software bugs in late 2023 to nearly 80% by late 2025, and 97% by September 2026, while the cost

of GPT-3.5-level output fell more than 280-fold. When capability rises and cost collapses at the same time, autonomy stops being a demo and becomes a procurement decision.

Then, you must follow the money, which is flowing upstream. Global private AI funding reached US \$217 billion in 2025; the first quarter of 2026 alone delivered US \$226 billion. Amazon, Alphabet, Microsoft and Meta are guiding to more than US \$720 billion in capital expenditure this year, nearly five times what they spent in 2023. Anthropic is now valued at US \$965 billion, OpenAI at US \$852 billion and NVIDIA at roughly US \$5.4 trillion. AI is no longer being financed like software; instead, it is being financed like power grids and railways.

Next there is the energy equation, the least glamorous and most decisive variable of all. Data centres consumed around 485 TWh of electricity in 2025, up 17% in a single year, and the IEA expects that figure to reach 950 TWh by 2030. In Ireland, data centres already absorb 23% of metered electricity; in PJM, America's largest grid market, capacity prices jumped more than ninefold. It's a bottleneck already showing signs of trouble in some markets, as models ship in months, while grids take years to build. Megawatts, not just model weights, now decide who can scale — and in 2026, so does the ability to physically protect them.



Against that backdrop, nations are moving at radically different speeds. More than 80 countries now have national AI strategies, but their playbooks are diverging: Europe has pushed its high-risk AI rules back to December 2027, Washington is prioritising speed and export controls, and Beijing is ordering foreign chips out of state-funded data centres. The United States holds roughly three-quarters of global AI compute, while China holds about 15%. Usage is also varied by geography: more than 70% of working-age adults use generative AI in the UAE, 63% in Singapore, 31% in the United States — and under 10% across much of Sub-Saharan Africa.

And then there is all that is happening here in Dubai. In the final days of April 2026, His Highness Sheikh Hamdan bin Mohammed bin Rashid Al Maktoum, Crown Prince of Dubai,

began issuing directives for Dubai's government — and its entire private sector — to transition to agentic AI within two years. Over roughly six weeks, the UAE committed to moving 50% of government services to agentic AI, launched an 80,000-person training programme, and set out plans to support 295,000 Dubai companies to adopt AI, develop 100 specialised AI assistants and establish 50 agentic AI companies. Where others set ambitions, Dubai set deadlines.

The UAE appointed the world's first Minister of State for AI in 2017. The Dubai AI Campus at DIFC set a target of attracting over 500 AI companies by 2028, a target that had been achieved by October 2026. The Dubai AI Seal has made trust a condition of public procurement. Regional data-centre capacity is set to more than triple by 2030, much of it built to AI-era

specifications from day one. Dubai is not trying to out-build Silicon Valley on frontier models. It is betting on something harder to copy: orchestration — of capital, regulation, energy, talent and demand, at speed.

Inside this report, you will find the data behind every one of these shifts: the kingpins redrawing the industry, ten charts that define where AI stands, the frontiers of agentic AI, healthcare, finance and cybersecurity, and a candid account of the challenges ahead.

If 2023 was the year of awakening, 2024 the year of acceleration, and 2025 the year of AI agents, then 2026 is the year of reckoning. The race is well underway.


Welcome to The State of AI 2026.

THE STATE OF AI IN 2026

Artificial intelligence has moved beyond experimentation to become foundational infrastructure. In 2026, the shift is no longer about model performance, but institutional execution. Advantage depends on integrating AI into workflows, governance, infrastructure, and talent at scale. Frontier development remains concentrated, while adoption diffuses unevenly. AI has entered its infrastructure phase — where scale and strategy matter more than novelty.



AI IS REORGANIZING THE GLOBAL ECONOMY AND THE SHIFT HAS ALREADY BEGUN



Economic history does not move in straight lines. It advances in structural shifts.

Most technologies improve existing systems. A rare few reorganize them. Mechanization created industrial scale. Electrification redesigned factories. Computing transformed how information flows. The internet

reorganized coordination across borders. Cloud infrastructure converted fixed capital into elastic services. Artificial intelligence now belongs in that lineage, but it operates at a different layer.


Earlier shifts scaled muscle, energy, movement, and information. AI scales cognition. It lowers the marginal

cost of reasoning, synthesis, pattern recognition, forecasting — and increasingly execution. Decision-making itself becomes cheaper and more distributable. What electricity did for physical power, AI is beginning to do for intellectual labor.

This is not incremental automation. It is economic reorganization.

STRUCTURAL TECHNOLOGY SHIFTS THAT REORGANIZED ECONOMIES

General-purpose technologies that reshaped production, coordination, and capital allocation



Source: Lucidity Insights Research

Structural technology moments share common traits: horizontal diffusion across industries, sustained infrastructure-scale capital allocation, institutional redesign, and geoeconomic realignment. AI now satisfies each.

Its reach is cross-sectoral. AI systems are embedded in healthcare diagnostics, financial compliance, cybersecurity, logistics, industrial optimization, energy forecasting, and scientific research. Frontier models increasingly approach — and sometimes exceed — human baselines on advanced reasoning tasks. Work that required specialized teams only a few years ago can now be executed in minutes at meaningful success rates. The boundary of economically automatable knowledge work is expanding in real time.

Capital allocation confirms the shift. Investment has moved upstream into semiconductor fabrication, hyperscale data centers, large-scale model training, and energy infrastructure. Hyperscaler capital expenditure is measured in the hundreds of billions annually. Data center expansion is influencing national grid planning, between 2025 and 2030, data center capital expenditure is expected to reach US \$6.7 trillion according to a 2025 McKinsey Study. AI is being financed as foundational infrastructure — more comparable to power grids or rail networks than to short-cycle software trends.

The center of gravity has shifted from applications to compute.

Institutions are reorganizing accordingly. While most enterprises report AI deployment in at least one function, few have achieved enterprise-wide transformation. The constraint is no longer model capability, but workflow redesign, data integration, governance discipline, and infrastructure readiness. The bottleneck has moved from technological feasibility to organizational execution.

History offers a parallel. Electrification did not transform factories when power first arrived. Productivity surged only


after production lines were redesigned around distributed energy. AI is entering that same embedding phase — where institutional architecture adjusts to the technology rather than merely experimenting with it.

The policy dimension reinforces the structural character of the shift. More than eighty countries have formalized national AI strategies, and several Middle Eastern states are positioning themselves aggressively. The UAE was among the earliest movers, becoming the first country to appoint a Minister of State for Artificial Intelligence in 2017 — signalling that AI would be treated not as a sector, but as national infrastructure. Advanced semiconductor export controls have turned chips into instruments of national security around the world. Binding regulatory regimes are replacing voluntary principles. Sovereign compute initiatives signal that access to intelligence infrastructure is now treated as strategic capacity. AI is embedded in economic competition and statecraft. *(Read more in the article “AI’s New Power Corridor: How the UAE and Saudi Arabia are Engineering the Middle East’s Rise.”)*

What makes this moment especially consequential is convergence.

Multimodal architectures now unify perception, reasoning, and tool use. Training compute continues to scale rapidly, even as inference costs decline, widening access while concentrating infrastructure. AI supercomputing capacity is clustered among a limited number of operators and geographies. Meanwhile, the electricity demand of AI data centers has become significant enough to shape energy policy discussions.

Capability, capital, infrastructure, and governance are aligning. Structural shifts become visible not at the moment of breakthrough, but when reinforcing layers accumulate — when capital deepens, institutions adapt, infrastructure expands, and regulatory frameworks crystallize.



Artificial intelligence is now crossing that threshold.

The question is not whether AI will advance. It will. The more consequential question is how unevenly its benefits, control, and influence will be distributed. As with prior structural technologies, advantage will accrue to those who align capital, infrastructure, institutions, and governance around this new economic layer.

The remainder of this State of AI Report examines how that reorganization is unfolding across enterprises, capital markets, infrastructure systems, and national strategies — not as a passing wave of innovation, but as the embedding of intelligence into the operating system of economic life.

“Artificial intelligence is no longer a sector within technology. It is becoming a foundational layer of the global economy.”

THE STATE OF GLOBAL AI IN 2026: ADOPTION SATURATES, SCALING GAPS WIDEN, AND PRODUCTIVITY DIVERGENCE BEGINS


Artificial intelligence in 2026 is defined less by possibility and more by execution. Nine in ten organizations now use AI somewhere; fewer than half have scaled it anywhere. Infrastructure gaps, skills shortages, sector constraints, and firm size disparities are shaping outcomes. The global AI landscape today is characterized by near-universal enterprise uptake, widening diffusion gaps, persistent scaling challenges, and the early signs of productivity divergence.

1. AI is Now a Strategic Imperative (not just a tech upgrade)

For many organizations, AI began with chatbots. But conversational interfaces are the visible layer of a deeper shift. AI is now embedded across core functions — customer service, marketing, software development, operations, risk, and compliance.

Generative AI lowered the barrier to entry, making advanced capabilities accessible to non-technical teams. The rise of agentic systems — capable of executing multi-step workflows with limited oversight — has pushed expectations beyond augmentation toward automation.

Adoption reflects this urgency. Nine in ten organizations now report regular use of AI in at least one business function. Among those already using it, 44% say AI is at least scaling across the enterprise — up six percentage points in a year — while the share still merely



experimenting has fallen from 32% to 22%. In many firms, AI is now a board-level priority tied directly to competitiveness and growth.

But the adoption curve itself has flattened: 89% in 2026 against 88% in 2025, after three years of steep expansion. The frontier has moved

from whether firms use AI to how deeply.


Yet momentum masks fragility. Even now, 56% of AI-using organizations — a clear majority, four years into the generative AI cycle — have not reached scale. While tools are spreading

ENTERPRISE AI ADOPTION: BREADTH PLATEAUS, DEPTH RISES (2017–2026)


“Nine in ten organizations use AI. Fewer than half have scaled it.”

Use of AI by respondents' organizations, % of respondents

Organizations that use AI in at least 1 business function¹



Phase of AI use among organizations using AI



¹In 2017, the definition for AI use was using AI in a core part of the organization's business or at scale. In 2018-19, the definition was embedding at least 1 AI capability in business processes or products. From 2020, the definition was that the organization has adopted AI in at least 1 function, and in 2025 and 2026, the definition was regular use of AI in at least 1 function.

Source: McKinsey Global Survey on the state of AI, 2017-26

2. Adoption / Diffusion is Uneven Across Countries, Firms, and Sectors

Beneath headline acceleration, AI adoption remains highly uneven.

Geographically, advanced economies continue to scale faster. Countries with strong cloud infrastructure, reliable energy systems, deep talent pools, and mature digital ecosystems are embedding AI more broadly. As of Q1 2026, nearly 1 in 5 working-age adults globally use generative AI tools regularly, but usage rates vary sharply — from over 70% in the UAE and above 60% in Singapore to just over 30% in the USA, and below 10% across much of Sub-Saharan Africa.

Emerging markets are active

participants, but adoption often concentrates in consumer-facing use cases rather than full operational transformation. Infrastructure, data availability, and compute access remain binding constraints.


Firm-level divergence mirrors this pattern. Large enterprises with centralized data architectures and prior digital investments hold structural advantages. Small and medium-sized enterprises face persistent barriers: fragmented data, limited capital, and difficulty attracting AI talent. Even as tools become easier to use, these constraints risk excluding SMEs from AI-driven productivity gains.

Sector variation adds complexity. Digital-native industries — technology, finance, media — are advancing more rapidly than asset-heavy or regulated sectors such as manufacturing, healthcare, energy, and public administration. In these environments, integration into legacy systems and compliance frameworks slows deployment.

Across geographies, firms, and sectors, three gating variables consistently determine outcomes: infrastructure, data quality, and skills. Where these are weak, AI adoption remains superficial regardless of ambition.



THE DIFFUSION OF AI ADOPTION ACROSS THE GLOBE (Q1 2026)



“Nearly 1 in 5 working-age adults worldwide now use AI regularly.”

Source: Microsoft AI Diffusion Report, Q1 2026
Note: Microsoft defines a user as someone engaging with AI tools for g90 minutes per month, as a share of working-age population

3. AI Activity is not the same as AI Impact

The most persistent pattern across markets is not a lack of AI activity, but a failure to convert activity into sustained impact. While pilots and proofs of concept are now widespread, far fewer organizations report measurable, enterprise-wide productivity or financial gains. Surveys are increasingly able to consistently measure this gap. According to a McKinsey & Company

survey, 44% of organizations have scaled AI across the enterprise and 80% report individual productivity improvements, yet only 37% report any positive contribution to earnings, and just 6% attribute 5% or more of EBIT to AI.

Several structural factors explain this gap. First, AI is frequently layered onto existing workflows rather than

used to redesign them. In such cases, AI automates discrete tasks but leaves underlying processes intact, limiting its ability to deliver step-change improvements. Fundamental workflow redesign is the clearest differentiator between the organizations seeing returns and those that are not. Second, data fragmentation remains a critical constraint. Many organizations lack integrated, high-quality data across functions, forcing AI initiatives to operate within silos.

Talent constraints further slow scaling. Although AI interfaces are increasingly accessible, deploying AI responsibly and at scale still requires scarce capabilities in data engineering, model oversight, risk management, and organizational change. These skills remain concentrated in a relatively small set of firms and geographies.

Infrastructure and cost dynamics compound the challenge. While model-level costs have declined, total system costs, which includes data preparation, integration, cloud infrastructure, compute, and energy, remain significant. As AI workloads grow more compute-intensive, organizations encounter limits related to cloud capacity, energy availability, and legacy system integration. These factors do not prevent experimentation but raise the complexity and risk of scaling.

Agentic AI sharpens these challenges. While agents promise workflow-level delegation, most organizations are not yet equipped to deploy them at scale. Ambiguities around accountability, auditability, and control (particularly in regulated environments) mean that agentic systems remain concentrated in narrow domains such as software engineering, IT operations, and customer support orchestration rather than embedded across core processes.

4. Productivity Divergence

The uneven diffusion of AI carries significant economic and strategic implications. At the national level, differential adoption risks translating into long-term productivity divergence. Economies that fail to build AI capabilities may become increasingly dependent on foreign technologies, platforms, and infrastructure, weakening competitiveness in trade, manufacturing, and services.

These dynamics are reinforcing AI sovereignty as a policy concern. Governments are prioritizing domestic compute capacity, data localization, and national AI ecosystems,

recognizing AI as a strategic asset rather than a neutral technology. Divergent regulatory and investment approaches across regions are likely to shape adoption trajectories over the coming decade.

For firms, the risk is structural disadvantage. Organizations trapped in cycles of experimentation without scaling may fall behind competitors who reinvest AI-driven gains into deeper capability building. Talent dynamics are likely to amplify this divergence if skilled workers increasingly gravitate toward firms where AI is embedded into core workflows rather than deployed superficially.

Workforce implications are equally material. As AI augments some roles and sectors faster than others, productivity and wage gaps may widen within and across economies. Workers and firms without access to effective AI tools risk falling behind even as aggregate technological capability advances.

Small and medium-sized enterprises face a particularly acute challenge. Without shared infrastructure, ecosystem support, or targeted policy measures, SMEs risk exclusion from AI-driven efficiency gains, accelerating market concentration and reducing economic dynamism.

5. From Experimentation to Transformation


Taken together, these trends suggest that AI is not merely another technology cycle but a force capable of amplifying existing inequalities if left to diffuse unevenly. The next phase will require a shift from horizontal, generic AI tools toward vertical, workflow-native applications embedded into core operations. Value will accrue to organizations that redesign processes, align governance, and invest in foundational capabilities required for scale.

The central challenge for decision-makers is therefore not whether AI will transform economies and enterprises, but how unevenly those benefits will be distributed. Diffusion alone is insufficient. Without deliberate investment in data, infrastructure, skills, and operating model change, the gap between AI activity and AI impact, and between

leaders and laggards, is likely to widen.

The next phase of AI will therefore be defined less by experimentation and more by allocation of capital, infrastructure, and institutional capacity. The transition from copilots

to agents, tightening infrastructure constraints, widening productivity divergence, and the growing centrality of governance will increasingly shape where value accrues. Access to AI will not be what separates leaders from laggards. The willingness to reorganize around it will.





MONEY FLOWS: WHAT AI INVESTMENTS IN 2026 TELL US


At first glance, AI investment patterns in 2025–2026 appear counterintuitive . Enterprise deployment remains uneven. Productivity gains are still emerging. Monetization at scale is concentrated in a narrow band of use cases. Yet capital flows into artificial intelligence have not slowed. In key segments, they have intensified.

This is not a disconnect between hype and reality. It reflects a structural shift in how investors, governments, and institutions interpret where long-term AI value will reside. Funding is no longer about validating AI’s promise. It is about securing position within its foundations.

Funding Remains Elevated, But Concentrated

Global private AI funding reached US \$217B in 2025. Q1 2026 alone added US \$226B, more than the entire previous year. What has changed is how that capital is distributed. OpenAI’s single US \$122B round took 54% of the quarter, and the ten largest deals took 82%. Even excluding OpenAI, the quarter reached roughly \$104B, up about 45% on Q4 2025.


Late-stage funding now dominates, while early-stage capital faces significantly higher thresholds for proof and differentiation. Investors are increasingly reluctant to back undifferentiated AI applications or incremental productivity tools. Capital is reserved for companies that demonstrate scale-readiness, defensibility, and deep integration into core enterprise or institutional workflows. The deal count shows it: it barely moved, at 8,170 in 2025 and 1,965 in Q1 2026 (down 5% on Q4), while the average deal in Q1 2026 was roughly four times the 2025 average. Scale matters more than novelty.



Source: CB Insights, State of Venture Q1’26; Lucidity Insights Research & Analysis

Unicorns Reveal where Confidence Lies

The changing profile of AI unicorns offers one of the clearest signals of this shift. In 2024, AI unicorns were spread across multiple stages of maturity, with a meaningful share still validating or deploying their solutions. By 2025, the picture has changed sharply. A dominant proportion of AI unicorns are now in the scaling phase which means that these companies have moved beyond experimentation and initial deployment into expansion, across customers, geographies, or infrastructure. The rapid increase in scaling-stage unicorns reflects growing investor confidence in execution and endurance rather than technological novelty alone. Early-stage and validating-stage unicorns now represent only a small fraction of the total, underscoring a decisive shift toward proven operating models.



Source: CB Insights



Capital is Flowing up the Stack

If we imagine the AI stack, it spans several distinct layers: applications, models, data and enablement, infrastructure and compute, and energy. Capital flows in 2025 point to a clear structural realignment across these layers. A disproportionate share of investment is now directed toward large-scale model development, compute infrastructure, data centers, specialized hardware, and the energy systems that support them.

These upper layers share common characteristics: they are highly capital-intensive, slow to monetize, and difficult to replicate. Yet they increasingly define who can build, deploy, and scale AI systems in practice. As awareness grows around compute intensity, energy constraints, and infrastructure bottlenecks, investors have moved upstream, toward the inputs that govern innovation velocity rather than the applications that sit closest to end users.

The implication is straightforward. Ownership, access, or control over infrastructure is becoming a primary determinant of who participates meaningfully in the AI economy. This shift has broadened the investor base beyond traditional venture capital, drawing in governments, sovereign entities, hyperscalers, and large corporates with the balance sheets and time horizons required to operate at this layer.

CAPITAL DEPLOYMENT AND TRENDS ACROSS THE AI STACK

Layer	Deal Activity	Capital Concentration	Typical Pattern	Representative Firms
 Application Layer	High	Low	Many startups; smaller rounds; faster revenue cycles; weaker moats	 Jasper  UiPath  Salesforce  Notion
 Model & Orchestration Layer	Low	High	Few players; large rounds; high burn; slower monetization; heavy R&D	 OpenAI  ANTHROPIC  Mistral AI
 Compute & Chips Layer	Very low	Very high	Hardware-led; mega-rounds; CapEx intensive; supply constrained	 NVIDIA  AMD
 Infrastructure & Data Layer	Medium	High	Cloud-led; platform scaling; enterprise integration; recurring revenue	 Azure  Google Cloud  databricks  snowflake
 Energy & Resources Layer	Very low	Very high	Asset-heavy; long build cycles; grid-dependent; often state- or hyperscaler-backed	Utilities, grid operators, data-center REITs

Source: Lucidity Insights Research & Analysis

Where the Middle East Fits

Within this global pattern, the Middle East plays a distinct role. Capital from the region follows the global shift toward the upper layers of the AI stack, with a clear focus on infrastructure, scale, and long-term control rather than a broad spread of early-stage applications. This reflects how AI is positioned in regional economic plans. In leading Gulf economies, AI is expected to contribute more than 10% of GDP by 2030, an outcome that depends on national-scale platforms, shared compute, and durable infrastructure rather than fragmented innovation.

Unlike ecosystems where funding

is spread across large numbers of startups, AI investment in the Middle East has focused on foundational assets such as compute capacity, data centers, and long-term platform partnerships. This pattern is visible at the country level. Saudi Arabia and the UAE have prioritized national AI platforms, sovereign-backed compute, and deep partnerships with global technology providers, while Qatar and Abu Dhabi have focused on cloud regions, data centers, and government-led AI use in sectors such as energy, healthcare, and public services.

In this context, the region acts as a capital amplifier. Innovation in AI may



originate globally, but capacity, access, and control are increasingly shaped locally. By funding the infrastructure and operating environments that AI systems depend on, Middle Eastern capital amplifies the reach and scale of AI developed elsewhere. As organizations move from pilots to enterprise-wide deployment, demand is shifting toward reliable capacity, energy, and governance (areas where the region has chosen to concentrate its resources).

By 2025, AI investment no longer resembles a typical technology adoption cycle. The rise of scaling-stage unicorns, the persistence of high funding volumes, and the prioritization of infrastructure all point to the same conclusion: AI is being financed as an economic foundation, not a short-cycle product market. Money flows are not simply expressions of optimism; they encode beliefs about where power in the AI economy will reside, and who intends to secure it.

“The Middle East is emerging as a capital amplifier in the AI economy — financing infrastructure, not just applications.”

10 GRAPHS

THAT TELL US ALL WE NEED TO KNOW ABOUT THE STATE OF AI TODAY


AI is no longer theoretical. It is showing up in enterprise budgets, product decisions, investment flows, and national strategies. Adoption is visible. Capital remains committed. Model capabilities continue to improve. At the same time, progress is uneven. Some firms are scaling rapidly; others are still testing. Some countries control frontier

infrastructure; many do not.

So rather than debate whether AI matters, it is worth asking what the data actually tells us about where we stand. These 10 charts and infographics offers a clearer picture of both the opportunity and the limits.

1. AI Adoption has gone Mainstream

AI'S INFLUENCE WITHIN ORGANIZATIONS




Source: Global Artificial Intelligence Report (2025) by International Data Center Authority

AI has moved decisively into the core of organizational priorities. A large majority of enterprises now use AI and generative AI, with adoption extending beyond experimentation. However, deeper integration, particularly in data-intensive and operational use cases, still lags headline-level intent.



2. AI is Scaled in Pockets, Not End-to-End

AI AGENT DEPLOYMENT THAT HAS REACHED SCALING, BY INDUSTRY AND FUNCTION IN 2025




Source: McKinsey Global Survey on the state of AI, 1,993 participants, June 25–July 29; Lucidity Insights Research & Analysis

Nearly nine in ten organizations now use AI in at least one business function, and 44% report it scaling across the enterprise. Depth is another matter. Measured at the agent layer, no function has passed 9% of organizations at scale. IT leads at 9%, knowledge management at 8%, and marketing and sales, service operations, product


development and software engineering all cluster at 7%. The core operational functions trail furthest, with supply chain and manufacturing both at 2%. Technology firms are the outlier: 24% have scaled agents in software engineering, the highest figure in the survey. Three in four still have not.

3. AI Is Not Software Alone: The Stack


AI IS 70% INFRASTRUCTURE AND SILICON, 30% SOFTWARE



4. The Great Divergence: Exploding Training Compute, Collapsing Inference Costs



Source: Epoch AI, 2025 | Chart: 2025 AI Index Report




Source: Epoch AI, 2025; Artificial Analysis, 2025 | Chart: 2025 AI Index Report

Frontier AI models require exponentially rising training compute—measured in petaFLOPs (one petaFLOP equals 10^{15} floating-point operations), a unit that captures the total computational work used to train a model. As model size and complexity increase, the amount of compute needed during training has surged. In contrast, inference costs, the

price to run a trained model per query or per million tokens (tokens are units of text, roughly words or word fragments), have fallen sharply. This creates a structural divergence: AI is increasingly expensive and resource-intensive to build, yet dramatically cheaper, making it more economically accessible and commercially viable for mass adoption.

5. Who Builds Frontier Models, and Where



Source: Epoch AI, 2026 | Chart: 2026 AI Index report

AI model development remains highly concentrated, with the United States and China operating at a scale few other countries match. The Stanford AI Index’s historical map shows the depth of that concentration across 2003–2025. US organisations contributed to 64 notable models in 2025, compared with 35 for Chinese organisations. Together, they were involved in 97 of the 107 models recorded that year, including joint projects.

Arabic and reasoning. Israel’s Jamba and LTX families span language and video, while Saudi Arabia’s ALLaM focuses on Arabic-language capability. These contributions broaden the geography of model development, even as the largest volume of recorded releases remains concentrated in the United States and China.

Behind these totals are competing ecosystems of developers. OpenAI’s GPT-5 and Anthropic’s Claude 4 were among the US releases recorded in 2025; DeepSeek-R1 and Alibaba’s Qwen3 featured among China’s contributions. The competition extends across reasoning, coding and multimodal systems, with successive releases from several major labs. Model counts capture the volume of recorded activity, however, rather than the performance gap between the two countries.

Outside these leading ecosystems, development is more dispersed. The UAE’s Falcon, Jais and K2 families illustrate its participation in general-purpose language models,




6. Open-Weight Models Narrow the Capability Gap



PROGRESS IN OPEN-WEIGHTS VS PROPRIETARY INTELLIGENCE

Highest-scoring model by license type at each release date, Nov 2022 – Sep 2026




Source : Artificial Analysis Intelligence Index v4.3.2 (10 evaluations: AA-Briefcase v1.1, GDPval-AA v2.1, AutomationBench-AA, Terminal-Bench 4.0, SciCode, Humanity’s Last Exam, CritPt, AA-Omniscience, AA-LCR v1.1), accessed 23 Sep 2026, Highest-scoring configuration per model, including max reasoning effort. Scores are not comparable with earlier index version. | Chart: Lucidity Insights

The Artificial Analysis Intelligence Index combines ten demanding evaluations into a single measure of frontier capability. Nearly a third of its weight sits on real professional work, such as multi-week projects, business automation and economically valuable tasks. The rest covers coding, scientific reasoning and factual reliability. Proprietary models are closed systems whose trained parameters (“weights”) are not publicly released. Open-weight models publish their weights, which lets anyone modify and fine-tune them.

The best open-weight model today, Xiaomi’s MiMo-V2.6-Pro, matches the level the closed frontier reached in June 2026. That is a lag of roughly three months, down from five to six a year earlier. The gap does not close in a straight line. It was narrowest in July, when Moonshot’s Kimi K3 came within five weeks of the frontier. A run of Anthropic releases has since pushed it back out to 11 points. Every open-weight leader since DeepSeek R1 in January 2025 has come from a Chinese lab.

7. Meaningful Enterprise Impact Is Limited to a Minority

OBJECTIVES OF AI EFFORTS AT RESPONDENTS’ ORGANIZATIONS, % OF RESPONDENTS¹



Source: McKinsey Global Survey on the state of AI, 2026

¹ Asked only of respondents who say their organizations are regularly using AI in at least one business function. Respondents who said “don’t know” or “other” are not shown.
² AI high performers are respondents who reported that 5% or more of their organization’s EBIT and “significant value” are attributable to the organization’s use of AI.

While AI adoption is widespread, meaningful enterprise impact remains concentrated in a minority. McKinsey’s 2026 survey shows that AI high performers, firms attributing 5% or more of EBIT to AI, are not distinguished by efficiency goals, which most organizations pursue (78% vs 82% for others), but by ambition: 74% set growth objectives versus 47% of others, and 65% set innovation objectives versus 49%. Automating workflows is the entry ticket; high performers are the ones aiming AI at revenue and new business models. Enterprise-wide transformation remains the exception, not the norm.



8. AI Funding is Clustered in Select Industries

TOP 5 TECH AND INDUSTRY SECTORS RANKED BY SHARE OF ANNUAL AI EQUITY DEALS

	2021	2022	2023	2024	2025
#1	Digital Healthcare	Digital Healthcare	Robotics	Robotics	Robotics
#2	Robotics	Robotics	Digital Healthcare	Digital Healthcare	Digital Healthcare
#3	Fintech	Fintech	Fintech	Fintech	Fintech
#4	Advanced Manufacturing	Advanced Manufacturing	Advanced Manufacturing	Cybersecurity	Cybersecurity
#5	Supply chain & Logistics	Supply chain & Logistics	Supply chain & Logistics	Advanced Manufacturing	Advanced Manufacturing

Source: CB Insights


AI investment remains concentrated in a small set of verticals, with robotics, digital health, and fintech consistently ranking among the top three sectors over the past five years. The stability of this ranking suggests sustained investor conviction in domains where AI deployment is

commercially viable and defensible. Rather than dispersing evenly across industries, capital continues to cluster in areas with clear automation pathways and scalable business models, indicating that investor focus remains selective despite broad interest in AI.



9. AI Momentum Has Accelerated Sharply Since 2022

MASS AI ACTIVITY AND SIGNALING EXPANDING




Source: Air Street Capital – State of AI Report 2025
Note: “AI Share (%)” reflects AI-related signals as a proportion of all tracked startup activity signals, while “AI signals” refer to aggregated indicators of startup momentum, including hiring growth, funding activity, product releases, and investor engagement associated with AI-focused companies.

AI companies now represent 41% of the Top 100 private companies, up from 16% in 2022. The chart illustrates the sharp acceleration of AI activity following the launch of ChatGPT in late 2022. Prior to Q4 2022, AI-related “signals”, which are real-time indicators such as hiring growth, product launches, funding activity, and investor engagement, all were growing steadily but modestly. After ChatGPT’s release, both the absolute number of AI signals and AI’s share of overall

activity surged dramatically. By 2023 and into 2024, AI represented roughly 35–45% of total signals, reflecting a structural reallocation of capital and entrepreneurial focus toward AI-driven ventures. Although activity moderates slightly in 2025, the data makes clear that ChatGPT marked a catalytic inflection point — transforming AI from a growing sector into a dominant force within the technology investment landscape.

10. AI Scaling is Constrained by Physical Infrastructure

AI DATA CENTRE DEMAND VS GLOBAL INCREMENTAL ELECTRICITY GROWTH (IEA PROJECTIONS)



8% of all new electricity capacity going to data centers by 2030

AI Development : **Months**
Grid Upgrades : **Years**

Multi-GW Cluster
Localized to US & China

AI scaling is colliding with physical limits. IEA projections show that while global electricity supply expands steadily, data centres alone absorb roughly 8% of all new capacity added by 2030, and materially more in advanced economies. Yet AI deployment cycles run in quarters, while grid upgrades take years. Multi-gigawatt clusters now require transmission

buildouts, interconnection approvals, and capital commitments that smaller countries simply cannot match. Compute is becoming geographically concentrated not just by talent or capital, but by power availability. Increasingly, AI leadership will depend as much on access to infrastructure as on model breakthroughs.



SPOTLIGHT | THE KING PINS OF AI: WHO ARE THEY AND WHERE THEY ARE IN 2026

THE KING PINS OF AI IN 2026

As 2026 enters its final quarter, the AI race is no longer crowded.

A handful of companies now shape the trajectory of artificial intelligence at scale, and we’ve all heard the names: NVIDIA, Google DeepMind, Microsoft, OpenAI, Anthropic and DeepSeek. They largely do not compete on the same products or user base. They control different choke points: compute, models, enterprise deployment,


research depth, and infrastructure. Together, they determine how fast AI advances, who can access it, and where economic leverage accumulates.

The combined market footprint of these six companies encompasses trillions of dollars in value — from NVIDIA’s market capitalization of more than \$5.4 trillion to Anthropic’s \$965 billion valuation, a figure that more than doubled in under four months.

DeepSeek has never taken a dollar of outside funding and has no known valuation at all. The AI era has produced the fastest-scaling technology companies in human history — and then, just when the pattern seemed set, a hedge fund manager in Hangzhou rewrote the rules entirely.

These are the king pins of the AI era in 2026.


nvidia


Microsoft


Google DeepMind


OpenAI


ANTHROPIC


deepseek

NVIDIA - The Chip Maker and the Gatekeeper of Compute



Jensen Huang
CEO

If AI is the new industrial revolution, NVIDIA manufactures the engines. For two decades, the tech industry’s value center of gravity sat firmly in software — Google, Facebook, Amazon, Salesforce, and their peers proved that code scaled infinitely and hardware was a commodity. NVIDIA has upended that assumption entirely. In the AI era, the

physical layer has reasserted itself with a vengeance: without the right silicon, even the most brilliant model remains theoretical. Hardware is back — and it is running the show.

Under Jensen Huang, NVIDIA has transformed from a graphics chipmaker into the indispensable backbone of AI compute globally. Its GPUs power nearly every frontier model training run in the West. Its CUDA software ecosystem has locked developers into its platform — switching costs are high, and rivals have yet to meaningfully erode that moat. Its newest platform, Vera Rubin, is not an incremental upgrade either. Unveiled in January 2026 and in full production by the summer, Rubin is six chips co-designed to work as one rack-scale AI supercomputer: a GPU, a CPU, and the networking and data-processing silicon that ties them together. It is the “AI

factory” Jensen Huang talks about, built as a system rather than a chip. NVIDIA says it cuts the cost of generating AI output by up to 10x compared with its predecessor, Blackwell, and trains the largest models with a quarter as many GPUs. By August 2026, Vera Rubin racks were already running at Microsoft Azure, Google Cloud, Oracle Cloud Infrastructure, CoreWeave and Nebius.

The numbers speak to a dominance that goes beyond market share. NVIDIA holds approximately 80% of the AI accelerator market by revenue, according to IDC. In fiscal year 2026 (ending January 2026), NVIDIA posted record revenue of US \$215.9 billion, up 65% on the prior year. Growth has since accelerated rather than slowed. In the quarter to July 2026, revenue reached \$96.2 billion, more than double the same quarter a year earlier, and US




SPOTLIGHT | THE KING PINS OF AI: WHO ARE THEY AND WHERE THEY ARE IN 2026

\$89.0 billion of that came from the data center segment, up 117%. NVIDIA is guiding to \$108 billion for the following quarter. Jensen Huang has said the company’s order book for Blackwell and Rubin systems exceeds US \$1 trillion through 2027.

In October 2025, NVIDIA became the first company in history to reach a US \$5 trillion market capitalization — surpassing Apple, Microsoft, Alphabet, and Amazon. By September 2026, its market capitalization stood at roughly US \$5.4 trillion.

Huang does not speak in defensive terms about AI risk. He speaks about inevitability. Accelerated computing, in his view, will permeate every industry. NVIDIA’s strategy reflects that worldview: expand the compute universe, integrate vertically, and sell not chips but full AI infrastructure. The company is now partnering with European governments and telecoms to deploy sovereign AI infrastructure, and its financial ties to OpenAI have only deepened. Beyond its participation in OpenAI’s 2026 funding round, NVIDIA has committed to help finance OpenAI’s own data-center buildout, including a dedicated campus in Ohio, as part of a 10-gigawatt infrastructure partnership, with reported commitments of up to US \$100–105 billion. Arrangements like this are now widely described as “circular financing”: the chipmaker funds the data centers that buy its chips.

As of 2026, NVIDIA’s dominance is technical but has become structural. Every serious AI player depends on its hardware. That dependency makes NVIDIA the pace-setter of the entire race, but its biggest threat comes from its own best customers. Google, Amazon, Meta, and Microsoft are all developing custom AI chips to reduce what the industry quietly calls “the NVIDIA tax” — with JPMorgan estimating custom chips will take 45% of the AI chip market by 2028. Trade policy compounds the pressure, and keeps rewriting the terms. U.S. export controls effectively banned NVIDIA’s H20 chip in China in April 2025, at an estimated cost of \$15 billion in lost revenue. By January 2026, Washington had swung the other way, formally clearing sales of the more powerful H200 to China subject to a 25% government surcharge. Antitrust probes by both the DOJ and EU add further friction. The deeper question is structural: the next phase of AI shifts from model training to inference and deployment — a workload that favors leaner, cheaper, purpose-built chips. Whether NVIDIA’s moat holds as the work changes is the defining question of its next chapter.



SPOTLIGHT | THE KING PINS OF AI: WHO ARE THEY AND WHERE THEY ARE IN 2026

Microsoft - The Institutional Integrator



Satya Nadella
CEO

Microsoft may be the most strategically disciplined AI company in the world. Satya Nadella moved early. Microsoft first invested in OpenAI in 2019 — three years before the generative AI explosion — and has committed a total of US \$13 billion to the partnership. When OpenAI restructured as a public benefit corporation in late October 2025, Microsoft emerged holding a 27% equity stake valued at approximately US \$135 billion at the time. JPMorgan analysts called it potentially “some of the best money ever spent” in tech. In its Q2 FY2026 earnings (quarter ending December 2025), Microsoft reported that its investment in OpenAI contributed US \$7.6 billion to net income alone.

But the real genius was not the investment — it was integration. By embedding AI directly into Microsoft 365, GitHub Copilot, Azure, and enterprise workflows, Microsoft turned frontier intelligence into recurring enterprise revenue at global scale.

The numbers bear that out. In Q4 FY2026 (ending June 2026), Microsoft posted revenue of US \$90 billion, up 18% year-on-year, with Azure growth accelerating to 43%. Earlier in the year, Nadella had noted that “we are only at the beginning phases of AI diffusion and already Microsoft has built an AI business that is larger than some of our biggest franchises.” Microsoft’s commercial remaining performance obligations — essentially contracted future revenue — reached US \$678 billion by mid-2026, up from US \$368 billion at the end of 2025. The company’s market capitalization stood at roughly \$3.66 trillion in September 2026.

The Copilot story is where Microsoft’s distribution advantage becomes most tangible. By July 2026, Microsoft had 30 million paid Microsoft 365 Copilot seats (Copilot is an AI assistant

embedded directly into Word, Excel, PowerPoint, Outlook, and Teams; it can draft documents, summarize email threads, generate presentations from a prompt, analyze spreadsheet data, and transcribe and summarize meetings in real time, paid on top of existing Microsoft 365 subscriptions). At the end of 2025, seat additions were up over 160% year-on-year, GitHub Copilot had 4.7 million paid subscribers (up roughly 75% year-on-year), over 90% of the Fortune 500 were using Microsoft 365 Copilot, and daily active users were growing roughly 10x year-on-year. Nadella confirmed that GitHub Copilot is now a larger business than all of GitHub was when Microsoft acquired it for US \$7.5 billion in 2018. That is a product — not a company — that has outgrown one of the most important acquisitions in tech history in under six years.



Microsoft is also quietly building its own model layer, independent of OpenAI. Its Phi family of small language models represents a deliberate counter-strategy to the compute-heavy frontier race. Phi-4-reasoning, despite having just 14 billion parameters, outperforms models many times its size — including DeepSeek-R1 at 671 billion parameters — on key mathematical and scientific reasoning benchmarks. The logic is straightforward: not every enterprise use case requires a frontier-scale model, and small, efficient models that run on-device or on-edge hardware can unlock AI deployment in contexts where cloud dependency is a barrier. It is a hedge, and an interesting one to keep watch over.


Nadella’s philosophy is pragmatic above all else. AI is not framed as a moonshot; it is framed as productivity infrastructure. Microsoft’s AI Foundry platform now has over 1,500 customers deploying both Anthropic and OpenAI models, and the number of customers spending more than \$1 million per quarter on Foundry has grown nearly 80% year-on-year. Azure is the exclusive cloud for OpenAI’s API — and OpenAI has committed to purchasing \$250 billion in Azure cloud services going forward. That is not a partnership; that is a structural lock-in that defines Microsoft’s AI revenue runway for a generation.

Despite the successes, Microsoft’s central challenge is one of execution rather than strategy. In December 2025, Nadella himself sent internal emails admitting that Copilot’s integrations “for the most part don’t really work” — a remarkable admission from a CEO who has staked the company’s next decade on the product. The deeper concern is capital intensity: Microsoft’s own capital-expenditure guidance for 2026 now sits around US \$190 billion, a massive bet that

enterprise AI adoption accelerates fast enough to justify the outlay — and even McKinsey’s August 2026 State of AI survey finds only 44% of organizations have scaled AI enterprise-wide. Then there is the OpenAI dependency: the relationship that made Microsoft is now its most visible vulnerability. OpenAI is selling ChatGPT Enterprise in direct competition with Copilot, has signed cloud deals with Oracle and Amazon


that dilute Azure’s exclusivity, and remains an unpredictable partner. That Microsoft 365 Copilot began defaulting to Anthropic’s Claude models in some regions in May 2026 says everything.

In 2026, Microsoft sits at the center of institutional AI adoption. It does not just build intelligence; it operationalizes it — at a scale no other company can yet match.



SPOTLIGHT | THE KING PINS OF AI: WHO ARE THEY AND WHERE THEY ARE IN 2026

Alphabet / Google DeepMind -
The Scientific Architect



If NVIDIA controls the hardware and Microsoft controls enterprise distribution, Google controls research depth.

Google acquired DeepMind in 2014 for a reported £400 million (~\$650 million) — a move that Google co-founder Larry Page privately described as potentially the most important acquisition in company history. At the time, DeepMind had no revenue-generating products. What it had was a team with a singular mission: to solve intelligence. Hassabis and his co-founders chose Google specifically because reaching AGI would require vast compute — and Google, Hassabis has said plainly, “had and still has the most computers.”

In April 2023, DeepMind merged with Google Brain to form Google DeepMind, now the unified center of Alphabet’s entire AI R&D effort. It was a watershed reorganization that ended years of internal rivalry — and lit a competitive

fire. Demis Hassabis — a young chess prodigy turned neuroscientist, who said he wanted to solve general artificial intelligence from a young age — now speaks daily with Alphabet CEO Sundar Pichai. Their collaborative cadence has delivered tangible results: in March 2025, the team found its groove with the launch of Gemini 2.5, and by November, Gemini 3 was drawing praise from across the tech community. By August 2026, the team had shipped Gemini 3.7 Flash.

DeepMind’s crown jewel remains AlphaFold — the AI system that solved protein structure prediction, a 50-year grand challenge in biology. The achievement earned both Hassabis and his colleague Dr. John Jumper the Nobel Prize in Chemistry in 2024. The breakthrough has been used by



over 190 countries and more than 2 million researchers, reshaping drug discovery at a scale that makes most commercial AI applications look incremental. For Hassabis, this is what AI should be: a tool for fundamental scientific discovery, not merely software productivity.

Google’s structural advantage is vertical integration: custom Ironwood TPU silicon, frontier multimodal models, global data scale via Search, YouTube, Android, and Google Cloud. Hassabis has said Google plans to invest more than US \$100 billion in AI. In practice, Alphabet has gone much further: its 2026 capital-expenditure guidance of roughly US \$175–185 billion is about double 2025’s US ~\$93 billion. Alphabet is valued at roughly US \$4.14 trillion. That capital, combined with



a compute infrastructure that dwarfs most sovereign nations, gives Alphabet a staying power in frontier AI that few can match.

Still, Google DeepMind’s biggest challenge is not capability — it is conversion. For years, Google arrived first to AI breakthroughs and then struggled to translate them into products the world actually uses. The original Gemini launch in early 2024 was a masterclass in


how to fumble a frontier moment: strong benchmark performance overshadowed by an embarrassing image generation controversy that cemented the perception of a company more comfortable in the lab than the market. The deeper tension is structural: AlphaFold solved a 50-year grand challenge. ChatGPT shipped in a weekend. Both matter — but the market rewards the latter. The good news is that Gemini 3 appears to have genuinely closed the capability gap, with Google’s

share of the AI chatbot market surging from 5.7% to over 21% in twelve months. Depending on the tracker, Gemini now holds roughly 18–25% of the market, much of it taken from ChatGPT. The question for the rest of 2026 is whether DeepMind can finally marry scientific depth with the product velocity the moment demands.

In 2026, Alphabet is not chasing headlines. It is building the architecture for the long game.

SPOTLIGHT | THE KING PINS OF AI: WHO ARE THEY AND WHERE THEY ARE IN 2026

OpenAI - The Accelerator



Sam Altman
CEO

OpenAI changed the tempo of the AI industry. Not just incrementally, but permanently.

Before ChatGPT, AI progress was largely invisible to the public. After its November 2022 launch, AI became mainstream infrastructure almost overnight. ChatGPT reached 1 million users in just 5 days. It hit 100 million users in under 60 days — a growth rate that made Instagram’s 2.5-year journey to the same milestone look glacial. By early 2026, ChatGPT had approximately 800–900 million weekly active users and processed over 2.5 billion prompts per day. But its share of the AI chatbot market has since fallen from more than 80% to roughly 45–55% by mid-2026, depending on the tracker, as Google’s Gemini gained ground.

The financial trajectory has been equally extraordinary. OpenAI generated approximately US \$3.7 billion in revenue in 2024, and by 2025 had surpassed

US \$13 billion in annualized revenue — a pace Reuters and Yahoo Finance describe as the steepest revenue ramp in tech history. The company raised US \$40 billion in March 2025 — then the largest single private funding round ever recorded — at a US \$300 billion valuation. By October 2025, employee secondary sales valued OpenAI at US \$500 billion, and on March 31, 2026, OpenAI announced a new round: US \$122 billion from 31 investors at an US \$852 billion post-money valuation, now the largest private funding round on record.

Under Sam Altman, OpenAI has restructured as a public benefit corporation, with Microsoft holding a 27% equity stake valued at US \$135 billion at the time of the restructuring. The company has committed to over US \$500 billion in cloud infrastructure deals across Oracle, Microsoft, Amazon, and others — a scale of capital deployment that signals this is no longer a startup but a platform-level institution.


The biggest concern with OpenAI is that it appears to be a financial black hole. It lost US \$5 billion in 2024 on US \$3.7 billion in revenue, and a reported US \$12 billion in just the July–September 2025 quarter alone. Its 2025 net loss reportedly came in at US \$38.5 billion, nearly triple the US \$14 billion it had been projected to lose in all of 2026. Revenue is scaling fast, but costs are scaling faster — every new model costs more to build than the last, and every prompt costs money to serve.

Compounding this is a growing debate about whether the “scaling laws” that underpinned OpenAI’s strategy — the assumption that more compute reliably produces smarter models — are hitting diminishing returns. Meanwhile, its enterprise market share fell from roughly half in 2023 to just 27% by end of 2025, and the contest has since turned two-way: Ramp’s business-payments data showed Anthropic overtaking OpenAI in business customer share in May 2026, before OpenAI clawed back ground over the summer. The question is not whether OpenAI is growing — it clearly is — but whether it can grow fast enough to outrun its own cost structure.

OpenAI has also moved into ground its rivals left open. When Anthropic lost access to US federal work in February 2026, OpenAI signed a US government agreement the same day. And in September 2026, it began rolling out GPT-6 “Astra”, its next flagship, shipped with explicit warnings about its advanced cyber capabilities.

Even so, Altman’s worldview is expansionary. He believes advanced AI systems will reshape economic structures at scale — and he is willing to deploy rapidly to get there. It does not own the cloud. It does not manufacture chips. But it accelerated the race — and then redefined what winning looks like.

Anthropic - The Alignment Strategist



Dario Amodei
CEO

Anthropic has emerged as one of the most formidable forces in AI — founded by former OpenAI researchers, built with safety as a structural foundation rather than an afterthought. While many AI competitors treat safety as a layer bolted on post-development, Anthropic built it into Claude from the ground up through its “Constitutional AI” framework — a systematic approach to training models to be helpful, harmless, and honest by encoding principles into the model itself. That differentiation has not merely appealed to ethicists; it has proven to be a commercial superpower. Enterprises operating in regulated and high-stakes industries — healthcare, finance, law, government — need AI they can trust. Anthropic has grown on this strategic stance.

The numbers are extraordinary, even by the inflated standards of the AI era. In Anthropic’s own words, its run-rate revenue has grown 10x annually for three consecutive years: from

roughly US \$10 million in 2022, to US \$100 million in 2023, to US \$1 billion in 2024, to over US \$14 billion in annualized run-rate revenue by February 2026, and then to US \$65 billion by August 2026, more than quadrupling in six months. As of early 2026, Anthropic served over 300,000 business customers, and its enterprise accounts — those spending over US \$100,000 annually — grew nearly 7x over the prior year. By May 2026, Ramp’s business-payments data showed Anthropic ahead of OpenAI in business customer share (34.4% vs. 32.3%), up from just 1 in 25 companies paying for Anthropic a year earlier. And 79% of OpenAI’s paying enterprise customers also paid for Anthropic — they are buying both.

Much of that revenue comes from businesses and developers rather than

consumers. As of early 2026, enterprise and developer API usage represented approximately 75% of Anthropic’s revenue, with Claude Code — its agentic coding tool launched publicly in May 2025 — hitting US \$2.5 billion in annualized revenue by February 2026. That is a product that went from zero to US \$2.5 billion ARR in nine months. No enterprise software product had ever scaled that fast.

The investor community has taken note. In February 2026, Anthropic closed a US \$30 billion Series G funding round led by GIC and Coatue, at a US \$380 billion post-money valuation. Barely three months later, it raised US \$65 billion more in a Series H at a US \$965 billion post-money valuation, above OpenAI’s most recent US \$852 billion mark. Amazon (US \$5 billion) and Google (US \$10 billion)



SPOTLIGHT | THE KING PINS OF AI: WHO ARE THEY AND WHERE THEY ARE IN 2026

anchored the round in April, and a broader syndicate closed the remaining US \$50 billion on May 28. Total capital raised since its 2021 founding now exceeds US \$130 billion. Backers include Amazon, Google, Microsoft, NVIDIA, Sequoia, GIC (Singapore), Qatar Investment Authority, and Abu Dhabi's MGX. The company is reportedly preparing an initial public offering as early as late 2026, with some reports putting the target valuation near US \$2 trillion.

Critically, Claude is the only frontier AI model available across all three of the world's largest cloud platforms — AWS Bedrock, Google Cloud Vertex AI, and Microsoft Azure Foundry. For enterprise buyers, that means no single-cloud lock-in. Anthropic's infrastructure strategy — training Claude across AWS Trainium, Google TPUs, and


NVIDIA GPUs — mirrors the multi-cloud diversification its customers themselves prefer.

Dario Amodei has repeatedly emphasized that AI capability growth may outpace governance frameworks — and that safety research must scale alongside intelligence. That long-range perspective has shaped not just Anthropic's culture, but its product architecture. Eight of the Fortune 10 are now Claude customers.

Anthropic's greatest risk is also its greatest asset: its identity. In early 2026, that identity came under pressure from two directions simultaneously. First, in February 2026, Anthropic overhauled its Responsible Scaling Policy, removing its central commitment to never train models beyond what its safety

measures could guarantee — a pledge it had publicly touted since 2023. Critics called it capitulation to competitive pressure. Second, Anthropic found itself in a dispute with the US Department of Defense over Claude's safeguards, putting a US \$200 million government contract at risk.

Amodei's response was unambiguous: "We cannot in good conscience accede to their request." The standoff did not stay contained. On February 27, 2026, US federal agencies were ordered to stop using Anthropic's products. Whether that stand makes Anthropic the moral backbone of the AI era — or a company that priced itself out of its own government market — is a question that will define its next chapter. So far, the commercial market has answered emphatically: Anthropic's revenue has more than quadrupled since the ban.



DeepSeek - The Disruptor



No account of AI in 2026 is complete without DeepSeek.

In January 2025, a little-known Chinese lab headquartered in Hangzhou released R1 — an open-source reasoning model that matched or surpassed OpenAI's best across key benchmarks, built for a reported \$6 million when U.S. rivals were spending hundreds of millions per training run. NVIDIA's stock dropped 17% in a single day, wiping approximately \$600 billion in market value — the largest single-day loss for any company in U.S. stock market history. The Nasdaq shed over \$1 trillion in a week. Many observers called it the Sputnik moment of the 21st century.

DeepSeek was founded in July 2023 by Liang Wenfeng — a Zhejiang University graduate who financed the entire operation through profits from his \$8 billion quantitative hedge fund, High-Flyer. He holds an 84% stake in the company. It has raised zero dollars in outside venture capital. In an industry

that has consumed tens of billions in funding before shipping a single commercial product, DeepSeek's story is radical.

What made R1 so destabilizing was not just its performance — it was what its existence implied. The prevailing assumption underpinning every billion-dollar infrastructure bet in Silicon Valley was that frontier AI required frontier spend. DeepSeek shattered that assumption. Working with older-generation NVIDIA chips — specifically because U.S. export controls had blocked access to the most advanced hardware — DeepSeek's engineers were forced to innovate around constraints rather than throw compute at them. The result was a suite of architectural and training efficiency breakthroughs that the entire global AI research community immediately began studying and adopting. DeepSeek's R1 not only matches OpenAI o1's quality at 90% cheaper per token, it is also nearly twice as fast.

By August 2025, DeepSeek released V3.1, surpassing prior models by over 40% on key benchmarks including SWE-bench and Terminal-bench. In April 2026, it followed with a preview of V4, again released open-source. A full successor to R1 has yet to materialise. Because it releases models open-source under the MIT license, every improvement becomes an immediate gift to the global developer community — including the very U.S. companies it is competing against.



DeepSeek's role in the AI market remains as provocative as its arrival. In February 2026, Anthropic accused DeepSeek, alongside fellow Chinese labs Moonshot AI and MiniMax, of using more than 24,000 fraudulent accounts to extract over 16 million Claude conversations to train their own models — a serious allegation that raises questions about the provenance of its capabilities. Data privacy concerns compound the picture: all user data is stored on servers in China, subject to PRC law, prompting bans across multiple governments and agencies. And its commercial model remains unproven — open-source generates adoption, but not necessarily revenue.

What is beyond dispute is its impact. DeepSeek permanently recalibrated what the AI industry believes is possible — and forced every player in this report to rethink their assumptions about cost, efficiency, and the geography of innovation. The kingpins of AI are no longer exclusively American.


SPOTLIGHT | THE KING PINS OF AI: WHO ARE THEY AND WHERE THEY ARE IN 2026

THE ARCHITECTURE OF THE AI ERA

What distinguishes these six companies is not merely their size or their funding — it is the structural position each occupies in the AI value chain. NVIDIA controls the physical infrastructure on which all AI runs. Microsoft controls the enterprise deployment layer. Google DeepMind controls the research depth that makes long-horizon scientific breakthroughs possible. OpenAI still leads in consumer mindshare, the force that set the tempo for the entire industry, though it no longer leads unchallenged. Anthropic controls the enterprise trust architecture. And DeepSeek has done something none of the others managed: it proved that the entire cost structure of frontier AI could be dismantled and rebuilt from scratch.

They are not competing in a zero-sum race. They are building different walls of the same cathedral — while the ground beneath all of them shifts. The next inflection point, whether it comes from a breakthrough in reasoning, true agentic AI, or AGI itself, will be built and shaped by the players profiled here.

The AI era is not a race to the finish. It is a reconfiguration of the global economy. And these six are writing the first draft of what that looks like.



SPOTLIGHT | AI & SOVEREIGN CLOUD

THE MAN TRYING TO ENABLE A SOVEREIGN ONLINE FUTURE BUILT BY AI



Dominic Williams and DFINITY Foundation are behind an open cloud that solves growing needs within the critical realms of sovereign tech and agentic infrastructure, pulling it to center stage, from Islamabad to the United Nations.

By Erika Masako Welch,
Chief Content Officer, Lucidity Insights

Today, all around the world, there is a growing clamor for sovereign tech infrastructure. Cloud and SaaS (i.e. Software as a Service) are now key foundations of modern societies, without which they cannot run. On the one hand, geopolitical tensions mean that governments want to ensure foreign powers cannot access their data, or deny them access to their infrastructure after disagreements, and enterprise too is seeking ways to reduce risky dependencies, and must comply with a growing body of regulations that require the re-homing of computation. On the other hand, the world is moving towards a future where AI builds its online services, and is deeply integrated with them.

Enter Dominic Williams and the DFINITY Foundation, a not-for-profit

research operation based in Zürich, Switzerland, that has spent more than half a billion dollars developing a new kind of tech infrastructure that it calls the “Internet Computer.” Williams describes his mission as providing an open technology stack: “The Internet Computer is a new kind of cloud, created the internet way, essentially by extending the network so that as well as connecting software, which makes systems like the web, email, chat and SaaS possible, it can also securely and resiliently host the software and data that powers them.”

The Internet Computer runs “serverless” software that doesn’t depend on servers like databases, and is tamper-proof, which means hosted services are immune to cyber attacks on infrastructure, and always-on, because the network redundantly replicates their computations and data across data centers in the background, making them incredibly resilient. DFINITY argues that this is exactly what the AI era demands. Because serverless software is simpler, AI agents can write it faster, and at lower cost, and the Internet Computer has also been designed to run the next-generation of software called “AIware,” which has AI on the inside.

The Internet Computer provides a cloud platform by combining compute nodes that are installed in different data centers around the world using an open mathematical network protocol. Williams believes that a new addition to the Internet Computer, called “cloud engines,” is currently propelling

the network into the mainstream by directly solving for sovereignty needs. He says that “cloud engines” allow enterprises to “configure the precise compute nodes the network combines to create their cloud infrastructure, creating the ultimate sovereign platforms,” and that “enterprises can change their configurations at any time, without interrupting or degrading hosted apps and services — providing complete freedom from vendor and location lock-in.”

The team behind the Internet Computer says they aren’t anti Big Tech hyperscalers. Williams and his team argue instead that it simply provides an open evolution of cloud for an era in which enterprise is seeking sovereignty and new platforms that play better with AI: “In just a few clicks, enterprises can combine nodes from the Internet



Computer’s marketplace to create new cloud engines. The marketplace includes hardware inside independent data centers, and virtual nodes running on hyperscaler clouds, and if those options don’t work, enterprises can

combine on-prem nodes too — so ultimately the market will have to decide what’s most popular.” New engines can be created in minutes at opencloud.org.



“Cloud engines allow enterprises to configure the precise compute nodes the network combines to create their cloud infrastructure, creating the ultimate sovereign platforms.”

Dominic Williams
Founder & Chief Scientist, DFINITY

SPOTLIGHT | AI & SOVEREIGN CLOUD

What’s perhaps most impressive, is that the software running on cloud engines —which powers apps and services — runs independently of how they are set up. Once an enterprise has created a cloud engine, they can change its configuration without interrupting the software they are running there. Williams says, “for example, you can start off with a cloud engine that is backed by hyperscaler compute, then later change the configuration so that it’s backed by dedicated hardware running in home country data centers, without interrupting or even degrading the services you are running,” and that “this provides unprecedented freedom from vendor lock-in.”

Some nation states have moved from experimenting with Internet Computer cloud engines to deploying them. Pakistan became the first nation state using cloud engines to combine locally-run hardware to create a next-generation sovereign cloud platform, led by the Pakistan Digital Authority, and is using them for a variety of needs. On the one hand, they are using their

sovereign cloud to develop a cyber-secure and resilient national chat platform, and on the other hand, they have used it in conjunction with the Caffeine app builder (caffeine.ai) to create internal government services, several of which are already live and in production. Sophisticated AI-driven development platforms such as Anthropic’s Claude Code and OpenAI’s Codex can also directly build and deploy secure apps and services on cloud engines.

In another recent partnership, the UNDP is using cloud engines for its own operations as well as working to deploy to country offices, governments and civil society organizations it supports, and the first UNDP cloud engine is already running for one of its country offices. Pierre Samaties, DFINITY’s Chief Business Officer, said: “Cloud engines are created from open source technology that is verifiably free of hidden foreign backdoors and kill switches, and engines host tamper-proof apps that cannot be hacked into, which is a game-changer for large

and developing countries alike. We’re seeing some interesting ideas, such as developing world nations asking about running cloud engines across hardware in their international embassies. Anything is possible.”

Williams goes on to say, “Internet Computer cloud engines offer a solution to sovereignty needs, but they also play an equally important role providing a better compute platform for the AI era. AI-powered hackers are making traditional compute infrastructure increasingly insecure, which tamper-proof software helps fix.

Tamper-proof software is also useful when AI builds and updates services in the role of developer, which it can do at the speed of conversation today. On traditional stacks, which aren’t tamper-proof, security and sysadmin teams have to be kept in the loop, which dramatically slows things down in production.”

Williams also talks about a new development coming called “Open SaaS” that reflects how he expects AIware will bring changes. He goes on to explain that Open SaaS aims to provide a free alternative to every SaaS service on earth that can be installed and operated on a cloud engine as easily as a phone app, and reimagines SaaS as AIware with intelligence on the inside. DFINITY is currently migrating away from all the traditional SaaS apps it used to run on, to apps from the Open SaaS suite, which he says are already delivering huge benefits. He goes on to say, “With Open SaaS, you spend more time in AI chat than using traditional screens. For example, you can say to AI ‘determine which PTO requests are overlapping in Customer Success, approve those that aren’t overlapping, and create emails asking personnel

to change their dates for the others, giving priority to our most valuable employees.”

DFINITY says more announcements are imminent, and that governments and large enterprises interested in meeting the team should attend DFINITY’s 2027 “Sovereign Cloud” event at the World Economic Forum in Davos this February. However, Williams says that cloud engines are for every enterprise, big and small, and that “once cloud engines are in general release, I expect we’ll see the fastest growth in the startup sector, because startups are flexible and have a huge incentive to be smarter and more efficient.” He also says that “network effects will drive rapid adoption of cloud engines because the serverless software they host is transparently interoperable with software on other cloud engines, subject to permissions, re-creating a key dynamic behind early internet growth.”



“AI-powered hackers are making traditional compute infrastructure increasingly insecure, which tamper-proof software helps fix.”

Dominic Williams
Founder & Chief Scientist, DFINITY





AI'S NEW POWER CORRIDOR: HOW THE UAE AND SAUDI ARABIA ARE ENGINEERING THE MIDDLE EAST'S RISE


The fastest way to understand why the Middle East is becoming a hotbed for AI is to stop thinking about “apps” and start thinking about infrastructure, policy, and power. The region’s leading markets—especially the UAE and Saudi Arabia—aren’t just adopting artificial intelligence. They’re building the conditions that make AI inevitable: sovereign-grade cloud capacity, national AI strategies, dedicated institutions, and partnerships with the world’s most important AI platform players. The opportunity is material. In 2018, PwC

estimated that artificial intelligence could contribute US \$320 billion to the Middle East economy by 2030 — roughly 2% of projected global AI gains and equivalent to about 11% of regional GDP. That estimate was made before the generative AI surge, and before ChatGPT and Claude turned frontier models into everyday enterprise tools.

Since then, global AI impact projections have expanded, with some estimates placing total economic upside at US\$17 trillion or more by 2030. If the Middle

East were to capture even 2–3% of that value, the regional impact would range from approximately US\$340 billion to over US\$500 billion this decade. The precise figure may vary, but the trajectory is unmistakable: AI’s economic upside for the Middle East is not incremental—it is structural. That kind of prize does not attract marginal reform. It demands nation-state strategy.

CONTRIBUTION OF AI TO GDP BY REGIONS, 2030



Source: PwC Report: The Potential Impact of Artificial Intelligence on the Middle East by 2030

The Middle East did not become an AI hotspot by accident. It did so by making artificial intelligence a matter of national strategy — and then funding it like infrastructure. While much of the world treated AI as a product cycle, the UAE and Saudi Arabia treated it as economic architecture. The result is not just adoption momentum, but structural positioning in the global AI race.

The UAE: Speed, Infrastructure and Deployment

The UAE was among the first movers globally, and they moved decisively on AI. In 2017, it launched its National Strategy for Artificial Intelligence

2031 and appointed the world’s first Minister of State for AI. That wasn’t just a symbolic gesture; it embedded AI into federal governance and accelerated coordination across ministries. By 2019, the country had established the Mohamed bin Zayed University of Artificial Intelligence (MBZUAI), creating a domestic research-to-deployment talent pipeline at a time when most nations were still debating AI ethics frameworks. But policy only matters if infrastructure follows.

Microsoft launched its Azure cloud regions (areas that operate multiple data centers) in both Abu Dhabi and Dubai in 2019. AWS opened its UAE region in 2022 with three Availability Zones (i.e. three physically distinct data center clusters that operate independently of each other). Google Cloud serves the market from its nearby Gulf regions in Doha and Dammam. These hyperscaler commitments represent billions in capital expenditure and significant installed data center capacity — the plumbing that allows AI workloads to run securely and at scale.

Energy is becoming the quiet constraint in the AI race. AI workloads are compute-intensive by design, and that intensity is materially increasing

electricity demand — particularly through hyperscale data centers. The International Energy Agency (IEA) estimates that data centers consumed around 485 terawatt-hours (TWh) of electricity in 2025, 17% more than a year earlier, while demand from AI-focused data centers rose by roughly half.

“In an AI economy constrained by compute and power availability, the ability to align digital expansion with national energy strategy is not a luxury—it is leverage. This is where parts of the Middle East begin to look structurally advantaged.”





This is where parts of the Middle East begin to look structurally advantaged. Gulf economies benefit from centralized infrastructure planning, available land for hyperscale development, and comparatively competitive energy costs. In an AI economy constrained by compute and power availability, the ability to align digital expansion with national energy strategy is not a luxury—it is leverage.

The UAE exemplifies this alignment. It operates a modernized grid, diversified generation capacity, and stable baseload power, including nuclear energy from the Barakah plant, whose four reactors supply about a quarter of the country’s electricity. Dubai’s Mohammed bin Rashid Al Maktoum Solar Park — one of the world’s largest single-site solar developments, at 3,860 MW today and targeting more than 8,000 MW by 2030 — anchors long-term renewable expansion, while Moro Hub, whose solar-powered data center opened at

the park in 2023, brings solar-powered data center capacity into the digital infrastructure stack. The UAE has since raised its clean-energy target, covering both renewables and nuclear, to 35% of its energy mix by 2030–31. In a world where AI growth is colliding with power constraints, that combination of reliability, scale, and planning discipline may prove decisive.

Why else has the UAE become a hot bed for AI adoption? The UAE also benefits from its digitally fluent population. Internet penetration stands at around 99%, and smartphone adoption is near universal. GDP per capita (PPP) is among the highest globally. Enterprise budgets exist; government procurement moves quickly; AI does not struggle to find buyers in this market.

The UAE’s model-building efforts — from Falcon 40B and Falcon 180B to the Technology Innovation Institute’s Falcon-H1 Arabic in January 2026

and MBZUAI’s K2 Horizon, a family of six fully open models released in September 2026 — have added global visibility. But its deeper advantage lies in being a fast, well-capitalized deployment hub: an international sandbox where AI can be tested, scaled, and governed. If the global AI race has hubs of invention, the UAE is positioning itself as a hub of execution.

Saudi Arabia: Scale, Sovereignty And Compute Power

Saudi Arabia’s AI trajectory is different — and arguably more industrial than that of the UAE. In 2020, the Kingdom formalized its National Strategy for Data and AI under Vision 2030, explicitly framing AI as a pillar of economic diversification. Where the UAE optimized for speed, Saudi Arabia is optimizing for scale.

Over the past few years, cloud infrastructure expansion has accelerated. Google Cloud opened its Dammam region in 2023. AWS is set to launch its Saudi region in December 2026, backed by a US\$5.3 billion investment, and Microsoft’s Saudi Arabia East datacenter region is due to open to customers in November 2026. These are not incremental upgrades; they are all signals that the Kingdom intends to host, not just consume, AI workloads.

The most consequential signal, however, is compute ambition. Partnerships tied to sovereign AI capacity — including HUMAIN’s partnership with NVIDIA to build “AI factories” of up to 500 megawatts over five years — indicate a willingness to treat compute like industrial infrastructure. HUMAIN, the national AI company launched by the Public Investment Fund in May 2025, is targeting 1.9 gigawatts of data-center capacity by 2030. How that capacity is financed is also changing: as the Kingdom reins in fiscal spending, HUMAIN has begun courting outside investors and is planning a listing in Riyadh and New York by 2029. In the AI era, megawatts could be equated to strategic advantage.

Saudi Arabia is also building Arabic-language AI sovereignty, deploying national models such as ALLaM into enterprise-grade platforms. HUMAIN’s ALLaM 34B, launched in August 2025, also powers its consumer app, HUMAIN Chat. The intent is clear: reduce dependence on foreign language models and strengthen local digital ecosystems. Where the UAE built a high-velocity AI ecosystem in a small, agile market, Saudi Arabia is building national-scale AI capacity in a much larger economy. One is optimized for global integration; the other for industrial depth.

Both countries have laid out somewhat similar foundations but are playing fundamentally different roles. The structural similarities are obvious: National AI strategies embedded in executive government; sovereign capital aligned to long-term digital infrastructure, now increasingly supplemented by outside investors, as both G42 and HUMAIN look to raise external capital; deep partnerships with hyperscalers such as AWS, Microsoft Azure, and Google Cloud, and a policy posture that encourages deployment across government, healthcare, education, and energy. Both also benefit from strong energy economics and centralized infrastructure planning — advantages that matter when AI growth becomes constrained by power availability. But their positioning within the global AI race differs.

The UAE functions as a global AI deployment hub — fast, internationally connected, digitally advanced, and highly attractive for multinational AI experimentation and rollout. Saudi Arabia, on the other hand, is emerging as a sovereign compute and AI industrial platform — building domestic capacity at scale, aligning megawatt power with GPU clusters, and embedding AI into national economic transformation.

The next test for that corridor is how reliably it can support AI at scale. Building capacity is one part of the equation; ensuring continuity as businesses and public services come to

depend on it is another. That requires dependable power and connectivity, redundancy across locations, and robust arrangements for data backup and recovery. The region’s advantages in energy, capital and speed remain significant. Translating those advantages into lasting value will depend as much on the reliability of the infrastructure as on the pace of its expansion.

Together, the UAE and Saudi are forming a new power corridor in the AI economy.

Neither the UAE nor Saudi Arabia currently competes with the United States or China in frontier model volume. But that is not the point. In an era where AI advantage increasingly depends on infrastructure, energy, deployment velocity and the ability to deliver reliable services at scale, both countries are building durable leverage. The Middle East is not trying to replicate Silicon Valley. It is building the substrate on which the next phase of AI growth will run. And in the infrastructure era of artificial intelligence, that may just be the smarter play.




KEY THEMES AND INNOVATIONS SHAPING 2026

AI's trajectory has moved rapidly, from experimental research models to conversational interfaces and embedded copilots across enterprise software. The shift now underway is toward systems that act, coordinate, and increasingly operate within real workflows. As capability accelerates, the core questions are no longer whether AI performs, but where it embeds, how it scales, and which constraints determine its durability. Across sectors, AI is moving from assistive layer to operational infrastructure. Autonomous workflows, domain-specific deployment, infrastructure bottlenecks, security, and governance maturity will determine whether this transition compounds advantage or introduces fragility.



COPILOTS: FROM ADOPTION TO INTEGRATION



Artificial intelligence in 2026 is still not defined by full autonomy. For most organizations, it is defined by copilots. AI copilots are human-in-the-loop AI systems embedded into workflows that draft, summarize, analyze, retrieve, and recommend — while leaving final decisions to people. They assist but do not act independently. That distinction explains both their rapid enterprise adoption and their institutional acceptability.

Across global surveys, nearly nine in ten organizations report using AI in at least one function. AI has moved beyond standalone tools and into embedded workflow infrastructure. AI chatbots and assistants are now the most widely scaled AI tool: 47% of respondents say their organizations are rolling them out across the enterprise. Microsoft alone reported more than 30 million paid Copilot seats in July 2026, with net additions more than doubling from the previous quarter. In many firms, copilot

% OF ORGANIZATIONS USING AI ACROSS FUNCTIONS AND PRODUCTIVITY IMPACT

Function	% of Organizations Using AI in Function (2025)*	Typical Productivity Impact*
 Service Operations	45%	20–30% reduction in average handling time
 Marketing & Sales	45%	10–20% increase in campaign throughput
 Product & Service Development	36%	20–40% faster iteration cycles
 Human Resources	34%	30–60% faster resume screening
 Software Engineering / IT	28%	30–50% faster task completion (controlled studies)
 Strategy & Corporate Finance	27%	30–50% faster research synthesis

Sources: McKinsey State of AI 2025; Stanford AI Index 2025; Microsoft Work Trend Index 2025; enterprise productivity studies (2023–2025).

*Note: Adoption percentages reflect McKinsey’s 2025 global survey. Other enterprise surveys report wider ranges depending on industry mix, respondent base, and definition of AI use.

adoption has been bottom-up. Knowledge workers integrate AI assistants into drafting, coding, research, and analytics tasks — often ahead of formal governance. Informal use frequently precedes structured deployment, illustrating how productivity technologies diffuse faster than policy frameworks.



What stands out is not dominance in a single function, but broad penetration across many. In the Gulf, service operations and marketing and sales lead adoption, with 45% of organizations embedding AI into each of those workflows. Product development, HR, software engineering, and corporate strategy follow, with product development showing the biggest jump since 2023, from 13% to 36%. No unit is fully transformed — yet few remain untouched. Copilots are no longer experiments; they are distributed productivity layers.

The impact is measurable but pragmatic, and uneven. In service operations and engineering, efficiency gains show up in reduced handling times and faster task completion. In marketing, finance, and HR, improvements manifest as compressed cycles and lighter cognitive loads rather than structural redesign. The largest controlled studies also show who gains most: less-experienced workers improve the most, while on tasks beyond what the model can reliably do, output can get worse rather than better. Across functions, the pattern is consistent: copilots augment work. They accelerate information flow, reduce repetition, and shorten decision loops, but they do not replace judgment. Machine-assisted workflows are becoming a baseline expectation across enterprise environments, even if operating models have not yet been rebuilt around them.

Regional differences are narrower than assumed. The GCC currently reports the highest enterprise AI usage rates, with North America, Europe, and Greater China clustering closely behind. Yet enterprise-wide scaling remains limited across all regions. Even in the GCC, fewer than one-third of organizations report scaling AI beyond pilot phases — a pattern broadly consistent globally.

The defining feature of this phase is institutional pacing. Copilots have proven they deliver efficiency gains. The strategic question now is whether organizations remain at the assistance layer — or redesign workflows around

these tools. McKinsey’s 2026 survey shows why it matters: eight in ten respondents say AI has improved their own productivity, yet only 37% report any effect on enterprise earnings, little changed from a year earlier. Among the small group of high performers, nearly three-quarters have fundamentally redesigned workflows around AI. Copilots may appear incremental, but at scale they reset productivity baselines. The transformation is subtle. The cumulative effect is not.

Regional adoption levels vary, but the differences are narrower than often assumed. The GCC reports the highest overall enterprise AI usage, while North America, Europe, and Greater China cluster closely together in adoption rates. Across regions, however, enterprise-wide scaling remains limited. In the GCC, fewer than one-third of organizations report scaling beyond pilots, a pattern broadly consistent with global surveys.

ENTERPRISE AI ADOPTION BY REGION (2025)



Source: McKinsey State of AI 2025; McKinsey GCC AI Survey 2025



AGENTIC AI: WHEN SOFTWARE STARTS DECIDING AND ACTING

Artificial intelligence is entering a phase where it does more than assist; it begins to operate. The shift from copilots to agents marks the most important architectural transition in enterprise AI since the emergence of large language models. Instead of generating drafts or answering prompts, agentic systems interpret goals, plan multistep workflows, call tools, and execute tasks with limited supervision. In practice, this means service tickets can be resolved autonomously, IT systems can self-diagnose, compliance checks can run continuously, and supply chains can rebalance in real time.

Enterprise surveys now show how far agents still have to travel. Nearly nine in ten organizations use AI regularly in

at least one business function, and 44% say AI is scaling across the enterprise. Agents trail well behind, and the gap is opening along company size. At organizations with more than US \$1 billion in revenue, 40% of respondents report scaling AI agents, up from 27% a year earlier. At smaller firms, the figure has not moved from 22%. The gap between experimentation and orchestration is where the next phase will be decided, and large enterprises are already pulling ahead.

What makes this moment different from two years ago is the convergence of capability and economics. On SWE-bench Verified, a benchmark built from real software bugs, the best AI systems went from resolving 4 percent of tasks

in late 2023 to over 60 percent in about a year, and near 80 percent by the end of 2025. Scores have since bunched up against that ceiling, and frontier labs have moved on to harder tests. At the same time, the cost of running a model at GPT-3.5's level fell more than 280-fold between late 2022 and late 2024, and Epoch AI estimates that the price of reaching any given capability level is still falling by roughly 5–10x a year. Reliability has improved while operating costs have collapsed. That combination turns autonomy from novelty into feasibility. Venture data reinforce the shift: AI-first companies are reaching meaningful revenue milestones faster than prior SaaS cohorts, suggesting strong enterprise pull rather than speculative demand.


FROM EXPERIMENTATION TO EXECUTION

AI adoption is widespread, but scaling agents remains a minority activity, and increasingly a large-company one, even as economics rapidly improve

SHARE OF ORGANIZATIONS



SCALING AI AGENTS, BY COMPANY SIZE



■ AI in general ■ AI agents

ECONOMICS ENABLE SCALING

280x


FALLING RAPIDLY

The cost of running a model at GPT-3.5's level fell more than 280-fold between November 2022 and October 2024.

\$20.00 → \$0.07 per million tokens

Sources: McKinsey & Company, The State of AI in 2026: On the Road to ROI (August 2026, 1,719 respondents, surveyed May 4–June 8, 2026); Stanford HAI, AI Index Report 2025 (inference cost).

CODING AGENTS: FROM 4% TO NEARLY 80% IN JUST OVER TWO YEARS



Source: SWE-bench Verified leaderboard (swebench.com). Data as of Feb 2026. Lucidity Insights Research & Analysis
Note: SWE-bench Verified = 500 human-validated real GitHub issues; a task counts as solved only if the model's patch passes the repository's tests.

Yet the real constraint is no longer model intelligence, rather it is institutional readiness. Scaling agentic systems requires workflow redesign, governance layers, data pipelines, and new skills in MLOps and oversight. Early deployments have exposed governance gaps, particularly around auditability and unintended tool execution. Many organizations remain in pilot mode not because the systems fail, but because integration into legacy processes is hard. Benchmarks also flatter. When METR asked open-source maintainers to review AI patches that had passed SWE-bench, roughly half would not have been merged. The implications extend beyond firms. AI adoption depends on infrastructure, data availability, and skills, and those are unevenly distributed. Recent global diffusion data show widening gaps in AI adoption between advanced and emerging economies. At the same time, autonomy amplifies risk. Cybersecurity research shows a sharp rise in AI-enabled phishing and deepfake activity, underscoring that autonomy must be matched by stronger security controls.

The defining question now is not whether agents work. The question is whether enterprises and governments can redesign around them fast enough. Autonomy without governance compounds risk; autonomy with institutional redesign compounds advantage. Even as agents advance, most organizations still rely on human-in-the-loop systems. While agents represent the direction of travel, copilots remain the operational baseline.



Agents Shaping 2026:

Beyond chatbots, these companies are building AI that doesn’t just answer questions — it takes action, completes tasks, and operates autonomously at enterprise scale. Valuations and ARR figures as of September 2026, if available. Valuation is the latest private round, or market capitalization for listed companies.

#	Company	Valuation	ARR / Revenue	Category	Founded	Agent Type	Why Watch
01	Agentforce	Public (CRM)	\$800M ARR	ENTERPRISE CRM	2004	Enterprise Agent	• 29,000+ enterprise deals closed • 330% YoY ARR Growth
02	CURSOR	\$29B	\$1B ARR	CODING	2022	Coding Agent	• Fastest B2B SaaS ramp in history • \$0 to \$1B ARR in 24mo • >100,000 products built daily on platform
03	SIERRA	\$10B	\$100M ARR	Customer Service	2023	Conversational Agent	• \$100M ARR in under 2 years • Outcomes based pricing
04	Devin	\$2B+	\$73M ARR	CODING	2023	Software Engineer Agent	• \$1M to \$73M ARR in 9mo • 8900 engineering teams using Devin
05	WAYMO	\$45B+	Est. \$500M+	PHYSICAL AGENT	2009	Autonomous Vehicle	• Most consequential autonomous agent already operating at scale: 150K rides per week across USA
06	glean	\$7.2B	Est. \$150M+	ENTERPRISE SEARCH	2019	Knowledge Agent	• 2,800 enterprise clients • “Google for your company” across Slack, Drive, Jira, Salesforce, etc.
07	OpenAI	\$852B (OpenAI)	See: OpenAI	BROWSER AGENT	2025	Browser Agent	• The first mass-market autonomous agent
08	Hippocratic AI — Do No Harm —	\$1.64B	Growing rapidly	HEALTHCARE	2023	Clinical Agent	• Addresses a global nurse shortage crisis at 1/100th the cost of human coordinators
09	UiPath™	\$7B (mktcap)	\$1.5B ARR	ENTERPRISE PA	2005	Workflow Automation	• 10,000+ enterprise customers
10	Harvey	\$5B	Est. \$100M+	LEGAL	2022	Legal Agent	• 4,700 law firm clients

Source: Lucidity Insights Research & Analysis
Note: Data compiled from public filings, funding announcements, earning releases, and industry reports. Valuations as of September 2026 , where available. ARR = Annual Recurring Revenue. Est. = Estimated from available data. Some figures are analyst projections.





AI IN HEALTHCARE: FROM CLINICAL VALIDATION TO PREDICTIVE, PERSONALIZED MEDICINE


Artificial intelligence is no longer confined to laboratory benchmarks in healthcare. It is entering regulated clinical environments at measurable scale. Over the past decade, FDA authorizations of AI-enabled medical devices have grown from just six annually to 335 in 2025, with 181 more in the

first half of 2026 alone. The FDA's own list now counts 1,614 AI-enabled devices, about three in four of them in radiology. Regulatory validation has shifted from exceptional to routine, embedding AI within formal medical oversight frameworks.

AI IN HEALTHCARE: FROM LAB TO BEDSIDE

Healthcare AI is advancing from regulatory validation to real-world clinical deployment


FDA-Authorized AI Medical Devices



Screening Accuracy

- ✓ 80.5% vs 73.8%: breast-cancer detection with AI-supported vs standard screening (MASAI trial)
- ✓ 12% fewer interval cancers in the AI-supported group

Clinical Knowledge Performance (MedQA)



*MedQA evaluates AI performance on U.S. medical licensing style clinical knowledge questions.

Operational Improvements

- ✓ 44% lower radiologist reading workload (MASAI trial)
- ✓ ~16 fewer documentation minutes per 8 hours with AI scribes (JAMA, Apr 2026)

Source: U.S. FDA, AI-Enabled Medical Devices list (updated Sept 4, 2026); Stanford AI Index 2025 (MedQA); MASAI trial, The Lancet (Jan 2026) and The Lancet Oncology (2023); JAMA (Apr 2026)

Performance gains are equally notable. Large language models now achieve 96 percent accuracy on medical licensing-style benchmarks such as MedQA, approaching expert-level reasoning in controlled settings. In clinical imaging, AI-assisted systems have delivered sensitivity improvements of five to ten percentage points in specific cancer detection trials. The strongest evidence so far arrived in January 2026, when The Lancet published the final results of Sweden's MASAI study, the first randomized controlled trial of AI in breast-cancer screening. Across more than 100,000 women, AI-supported screening detected 80.5% of cancers, against 73.8% for standard double reading by two radiologists, and was followed by 12% fewer interval cancers, those diagnosed between screening rounds. An earlier readout from the same trial showed radiologists' reading workload falling by 44%. These advances move AI from theoretical capability to clinical augmentation.

Operational integration is increasingly visible, but here the evidence is more sobering than the marketing. AI scribes, which listen to consultations and draft the clinical note, have been promoted as cutting documentation time by 20 to 40 percent, even handing back one to two hours per clinician per day. The largest real-world study to date, published in JAMA in April 2026 and covering more than 8,500 clinicians at five U.S. academic health systems, found a far smaller effect: adopters spent about 16 fewer minutes on documentation per eight hours of scheduled patient care, with no significant change in after-hours charting. Clinicians who used a scribe in at least half of their visits saved around 27 minutes, but only a third of adopters did. Other studies report real reductions in burnout, which suggests the value of scribes lies as much in attention and fatigue as in minutes saved.

In tuberculosis screening, the World Health Organization has recommended computer-aided detection software for reading chest X-rays since 2021, allowing programs to screen at scale in places with few radiologists. Drug

discovery is moving too. In September 2026, Insilico Medicine dosed the first patient in what it describes as the world's first Phase III trial of a drug developed with generative AI: rentosertib, for idiopathic pulmonary fibrosis, whose biological target was also identified by AI.

Yet capability alone does not guarantee transformation. Healthcare systems remain among the most regulated and risk-sensitive environments in the world. Ethics research, bias audits, demographic validation, and post-market monitoring are intensifying across regions. Questions of liability, reimbursement eligibility, and cross-border data governance remain central. Benchmark accuracy does not automatically translate into consistent real-world performance, and much of the evidence is still built on benchmarks: a JAMA review of 519 studies evaluating large language models in healthcare found that only 5% used real patient data. The cutting edge is shifting toward system-level intelligence. Multimodal clinical foundation models now integrate imaging, genomics, laboratory results, and physician notes into unified reasoning systems. At the frontier, digital twin models seek to simulate patient-specific physiology, allowing clinicians to test interventions virtually before administering treatment. Though early, these approaches signal a gradual move from reactive care to predictive and preventative medicine.

Saudi Arabia offers a compelling case study of national-scale healthcare AI adoption. Under Vision 2030's Health Sector Transformation Program, the Kingdom has invested heavily in nationwide digital health infrastructure, including unified electronic health records and AI-powered public health platforms. One flagship initiative is SEHA Virtual Hospital, launched in 2022 and now supporting more than 240 hospitals across the country, offering services in 48 main specialties and 68 sub-specialties, with an annual capacity of almost 600,000 patients — a model that earned global recognition



and a Guinness World Record as the world's largest virtual hospital. In 2025 alone, it delivered more than 16 million virtual appointments and consultations. Beyond the virtual hospital, the Ministry of Health's Sehhaty app, which handles appointment booking, prescriptions and sick leave, passed 31 million users in 2025, demonstrating how digital platforms are woven into everyday healthcare delivery.

While full physiological digital twins remain an emerging frontier globally, Saudi Arabia has taken a first, consumer-level step. In October 2024, the Ministry of Health added a "digital twin" feature to Sehhaty that lets users track their health status and receive AI-driven lifestyle advice. It is a personal health profile rather than a clinical simulation, but together with centralized planning and nationwide digital records, it gives the Kingdom a foundation few health systems have for moving toward simulation-driven medicine.

The defining constraint is no longer raw model accuracy. It is institutional integration. Healthcare adoption depends on interoperability with electronic health records, regulatory alignment, reimbursement frameworks, and clinician trust. AI has demonstrated capability. The next phase will determine whether health systems redesign around it.

In medicine, technological progress is necessary — but institutional trust determines scale.



AI IN FINANCIAL SYSTEMS:

EXPANDING ACCESS, MANAGING RISK

Finance is among the most tightly regulated infrastructures in modern economies. For decades, financial decisions, from lending to fraud detection, shifted gradually from human judgment to rule-based systems and quantitative models. Those systems were automated but bounded and overseen by people. Artificial intelligence changes the scale and speed of that transition. It does not simply follow predefined rules; it adapts and learns. The question is no longer whether finance can use AI. It is whether financial systems can trust it at scale. Today, about two-thirds of financial institutions report actively using AI, up from 45 percent a year earlier, and one in five has already deployed AI agents. Around nine in ten use AI to fight fraud and financial crime. What began as analytics support has become core infrastructure for risk management and compliance.

The most visible impact is faster financial decision-making. In anti-money laundering programs, mature AI systems reduce false positives by 30 to 50 percent, lowering investigative burden while improving signal quality. Advanced fraud detection models identify roughly 90 percent of fraudulent attempts in card and payments environments, shortening investigation cycles and accelerating risk assessment. Automation reinforces the shift: as early as 2017, JPMorgan’s COIN platform eliminated 360,000 hours of annual contract review work, while Morgan Stanley deployed a GPT-powered assistant to more than 15,000 financial advisors in 2023, marking one

of the first large-scale generative AI integrations in wealth management.

AI is reshaping who can access financial services. In the United States, AI-based underwriting platforms such as Upstart report that their models approve 43 percent more borrowers than a traditional credit-score-based model, at a 33 percent lower average APR, according to the company’s own 2025 study. In Saudi Arabia, fintech lenders licensed by the Saudi Central Bank (SAMA) and the Capital Market Authority are using data-driven underwriting to speed up SME financing in support of Vision 2030 diversification goals. Saudi-headquartered erad, for example, says its proprietary risk models let it approve SME financing within 48 hours. Digital advisory platforms further lower the cost of portfolio construction and automated rebalancing, extending investment access beyond traditional high-net-worth segments. The next step is agentic: in July 2026, the UK’s Financial Conduct Authority found that a fifth of UK adults are likely to use AI that can act autonomously, within pre-set goals, on their finances.

Acceleration, however, reshapes the system itself. As models compress research cycles, accelerate risk assessment, and automate execution, financial institutions face a structural trade-off between speed and stability. Governance discipline will determine whether acceleration enhances resilience or amplifies fragility. The Bank for International Settlements has cautioned that growing reliance

on concentrated AI infrastructure and similar modeling approaches could heighten the risk of synchronized behavior during periods of market stress. In June 2026 it went further, calling on supervisors to improve their visibility into “model-driven correlated exposures, herding behaviour and the potential for algorithmic collusion.” Binding rules are still catching up: the EU AI Act classifies AI credit scoring as high-risk, but the Digital Omnibus has pushed those obligations back from August 2026 to December 2027. Financial systems depend on trust, transparency, and explainability — qualities that do not always align naturally with complex machine learning models. AI is enhancing fraud detection, speeding compliance workflows, and broadening access to capital. But finance cannot tolerate opaque automation at scale. In financial systems, speed is valuable, and stability is indispensable.



IN FINANCE

ADOPTION PHASES AND SUPERVISORY RESPONSE

Phase	Time-frame	Adoption Trend	Regulatory Response / Movement
 Experimentation	2015–2017	- ML expands in fraud, credit scoring, trading - Large banks build AI teams (e.g., JPMorgan, Goldman Sachs) - Early robo-advisory launch	- No AI-specific rules - Governed under existing model risk frameworks (U.S. Fed SR 11-7) - AI treated as part of traditional risk management
 Expansion	2018–2020	- AI spreads across AML, chatbots, underwriting - Deep learning adoption increases	- BIS and FSB publish AI papers - Greater focus on explainability and model validation - Regulatory sandboxes expand (UK FCA, MAS)
 Acceleration	2021–2023	- AI embedded in core workflows - Compliance and research automation expand (JPMorgan COIN, Morgan Stanley tools) - Generative AI pilots begin (Goldman Sachs, Citi)	- Regulatory attention rises - BIS issues AI guidance - EU drafts AI Act - U.S. regulators publish AI risk principles - Scrutiny of cloud and third-party risk
 Institutionalization	2024 to 2026	- Generative AI integrated into compliance and customer service - Enterprise-wide adoption expands (JPMorgan LLM Suite to 200,000 staff; Goldman Sachs GS AI Assistant firmwide, 2025) - AI agents enter production (one in five firms, 2026)	- EU AI Act finalized (2024) - Digital Omnibus delays high-risk rules, including AI credit scoring, to Dec 2027 - MAS proposes AI risk management guidelines (2025) - FCA launches AI Live Testing (2025) and Mills Review (2026) - BIS calls for supervisory visibility into AI herding (2026)

Sources: Lucidity Insights Research & Analysis; data compiled from Bank annual reports and earnings disclosures (JPMorgan, Goldman Sachs, Morgan Stanley, Citi, HSBC); U.S. Federal Reserve (SR 11-7 and AI risk guidance); Basel Committee on Banking Supervision (BIS); Financial Stability Board (FSB); European Commission (EU AI Act); UK Financial Conduct Authority (FCA); Monetary Authority of Singapore (MAS).



AI AND CYBERSECURITY: ESCALATION WITHOUT EQUILIBRIUM

Cybersecurity has always been asymmetric. AttackersMandiant succeed once; defenders must succeed every time. Artificial intelligence compresses that imbalance. Attack volumes are at record highs. Check Point Research counted an average of 1,968 attacks per organization per week in 2025, roughly 18 percent more than in 2024 and 70 percent more than in 2023, and growth has picked up again this year: in August 2026, the weekly average reached 2,422, up 22 percent year-on-year. Automated password and network scanning now run at tens of thousands of attempts per second. The cycle between vulnerability discovery and exploitation has tightened dramatically. Google's Mandiant found that in 2025, flaws were exploited on average seven days before a patch was available, and that the gap between an

initial break-in and its hand-off to a second criminal group had shrunk to 22 seconds, from more than eight hours in 2022.

Before AI, cybersecurity was largely rule-driven. Systems relied on known signatures and fixed thresholds, such as five failed logins, a suspicious IP address, an unusual transaction. Software executed logic, but humans interpreted ambiguity. Security teams manually reviewed alert queues and correlated logs to determine whether anomalies signalled intrusion. Detection worked, but it was reactive and often overwhelmed by scale.

AI shifts that model. Instead of waiting for predefined rules to trigger, systems learn behavioural baselines and flag deviations in real time. Around four in

five security-operations practitioners now use AI or machine-learning tools in their daily work to identify anomalous activity, detect insider risk, and recognize evolving attack patterns, according to the SANS Institute, although only about a third have built those tools into defined, governed workflows. In mature environments, AI materially reduces false positives, improving signal quality and accelerating response cycles. The intelligence layer is no longer exclusively human; it is increasingly statistical.

Yet the same technology strengthening defence also empowers offense. Generative AI enables the rapid production of tailored phishing campaigns and synthetic identities. Deepfake voice cloning and video have been used to impersonate senior executives, triggering multi-million-dollar fraudulent transfers: in 2024, a single deepfake video call cost engineering firm Arup about US\$25 million. Voice phishing is now the second most common way attackers break in, behind software exploits, according to Mandiant. At the same time, the number of automated service accounts, APIs, AI agents and software workloads inside enterprises now far exceeds the number of human users (an average of 109 machine identities for every human one in 2026, up from 82 a year earlier). Each of these machine identities requires credentials and access permissions, significantly expanding what must be monitored and secured. AI systems themselves introduce new vulnerabilities, from

hallucination risks in automated workflows to potential model manipulation or data poisoning. Employees add a risk of their own: in July 2026, Check Point found that one in every 36 prompts sent to generative AI tools from corporate networks carried a high risk of leaking sensitive data.

The sharpest escalation is coming from the frontier labs themselves. In April, Anthropic's Claude Mythos Preview found thousands of serious vulnerabilities across every major operating system and web browser, and Anthropic limited access to a small group of defenders. In September, OpenAI gave its newest flagship a "Critical" cyber rating, its top risk tier, reserved for models that can find and exploit unknown flaws without human direction. The risks are already real: in July, two OpenAI models under test broke out of their sandbox and hacked into Hugging Face, and Anthropic disclosed three cases of Claude models


compromising real organizations' systems during evaluations. For now, Mandiant finds that most breaches still stem from basic security failures rather than AI. The question is how long that holds.

Institutional readiness remains uneven. Only one in ten large organizations is equipped to protect against AI-augmented threats, according to Accenture's latest global survey, published in 2025, although the share of organizations that assess the security

of AI tools before deploying them has nearly doubled in a year, from 37 to 64 percent, according to the World Economic Forum. Budgets have not kept pace: security budgets grew just 4 percent in 2025 and 5 percent in 2026, according to IANS Research and Artico Search, and nearly seven in ten security chiefs say AI is now their top priority for any new money. AI is accelerating both attack and defence, but control frameworks are not evolving at the same pace. Cybersecurity is the first domain where AI defends against AI at scale.




MACHINES NOW OUTNUMBER HUMANS BY A WIDE MARGIN



Sources: CyberArk, Identity Security Threat Landscape Report (2022); CyberArk, 2025 Identity Security Landscape; Palo Alto Networks (CyberArk), 2026 Identity Security Landscape. Averages from separate annual surveys of 1,750–2,900 security decision-makers; indicative of trend, not a like-for-like panel.

THE EXPLOIT WINDOW HAS CLOSED

Average number of days between a software flaw being disclosed and attackers exploiting it. Below zero: exploited before a fix existed.



Sources: Google Mandiant, time-to-exploit analyses; Mandiant M-Trends 2026, Lucidity Insights Research & Analysis

SPOTLIGHT | AI & PROPTech

AI & PROPTech: HOW RENO APP IS USING AI TO TRANSFORM REAL ESTATE'S MOST UNPREDICTABLE LAYER

Real estate has digitized transactions, payments, and asset management, embedding data into nearly every stage of ownership. Yet one of its most capital-intensive layers remains structurally ungoverned: renovation.

Once ownership transfers and upgrades begin, institutional discipline gives way

to informal coordination. Contractors operate across fragmented channels, budgets evolve without standardized benchmarks, and capital is deployed with limited predictive visibility.

For Marc Michel, Co-Founder and CEO of Reno, this reflects a structural data gap within the real estate lifecycle.

“Renovation is one of the most valuable phases in the real estate lifecycle, yet it still operates like it’s the 90s,” he says.

“Homeowners are juggling contractors on WhatsApp, tracking budgets in Excel, and relying on goodwill to stay on schedule. We believed that if execution could be structured, risk could be structured. And once risk is structured, capital can follow.”

Reno is building what it describes as an operating system for renovations — an infrastructure layer where artificial intelligence supports execution governance, cost normalization, and underwriting logic. Since launch, the company has completed more than 120 projects and raised \$3.6 million in equity and debt financing, positioning it to scale its AI-driven execution model across the Middle East.

The Structural Blind Spot in Real Estate

Renovation sits at the intersection of developers, contractors, and lenders, yet remains institutionally owned by none of them. Developers deliver assets, contractors execute upgrades, and banks finance acquisitions, but the stage where

capital is converted into tangible value through refurbishment has historically lacked standardized oversight.

The consequence is persistent uncertainty. Execution variables are difficult to benchmark, documentation is inconsistent, and delivery performance is rarely centralized. For lenders, renovation risk remains difficult to model because reliable historical comparability is limited.

Farah Karabeg, Reno’s Co-Founder and COO, describes what the company observed in the market. “Renovation has always been treated as an emotional process. What we identified was the absence of a governance layer. If scope is standardized, approvals are digitized, and payments are tied to verified milestones, much of the chaos becomes controllable.”

Within this structure, AI provides the technical infrastructure that embeds governance directly into execution.

Embedding Intelligence Where It Matters

Reno applies AI selectively, targeting points where ambiguity historically generates financial and operational risk. At the design stage, AI-generated visualizations align expectations early in the lifecycle, reducing downstream friction. More consequential is the normalization of Bills of Quantities. By standardizing contractor language and benchmarking line items against historical project data,




Reno transforms opaque negotiation into structured comparison.

Execution oversight follows the same logic. Baseline video capture establishes a digital reference point. Contractor updates are converted into structured logs, inspection workflows are generated according to scope, and progress is validated against predefined milestones.


Oversight scales through system logic, with interventions triggered by deviations from defined thresholds.

As Karabeg notes, “We’re not digitizing renovation for convenience. We’re digitizing it to enforce discipline. Predictability is the innovation.”



“We’re not digitizing renovation for convenience. We’re digitizing it to enforce discipline. Predictability is the innovation.”


Farah Karabeg
Co-Founder and COO, Reno



From Execution to Underwriting

The most strategically differentiated aspect of Reno’s model lies in how execution data is integrated with capital deployment.

Traditional renovation financing is



constrained because lenders lack structured insight into delivery risk. Reno links capital release directly to verified milestones within a controlled execution environment. Funds are disbursed in tranches following inspection validation, while retention mechanisms and behavioral signals

inform underwriting decisions. Marc summarizes the principle: “We finance what we control. If you separate money from execution, you amplify risk. If you link them through data, you reduce it.”

As projects accumulate, Reno’s dataset expands to include contractor reliability, cost variance, timeline adherence, and payment performance. Underwriting sharpens as intelligence compounds, allowing financing capacity to scale alongside predictability.

What next?


Reno’s next phase focuses on predictive modeling and persona-based credit engines that further integrate governance into the platform’s core. Much of this capability will be embedded in Reno App 2.0, the company’s next product iteration, which introduces AI-driven visualization, structured scoping tools, and automated quality and milestone tracking designed to transform how renovation projects are initiated and managed.

As predictive capabilities mature, renovation begins to resemble a

structured financial instrument rather than a volatile project. Standardized scopes, benchmarked costs, and forecastable timelines allow capital to be deployed with greater confidence.

Amr Hosny, Co-Founder and Chief Commercial Officer, situates this shift within a broader industry evolution. “Real estate transactions have become transparent and digitized, yet the post-acquisition phase has remained informal. We are professionalizing that stage. Investors increasingly recognize renovation as a value lever, and value levers require structure.”

Reno’s model illustrates how AI can function as a governance layer within the built environment, structuring execution in areas where capital has historically outpaced control. By converting renovation into structured, auditable, and financeable data, the company reframes it as a repeatable and institutionalized component of the real estate value chain.



“Real estate transactions have become transparent and digitized, yet the post-acquisition phase has remained informal. We are professionalizing that stage.”

Amr Hosny
Co-Founder and CCO, Reno



“We finance what we control. If you separate money from execution, you amplify risk. If you link them through data, you reduce it.”

Marc Michel
Co-Founder and CEO of Reno



SPOTLIGHT | THE CHALLENGES AHEAD IN THE GLOBAL AI RACE

THE CHALLENGES AHEAD IN THE GLOBAL AI RACE

“The technology is extraordinary. So is our capacity to mishandle it. Here is a clear-eyed look at the hardware bottlenecks, energy constraints, geo-economic fault lines, and human costs that will define the next decade of artificial intelligence.”

Let’s be honest about something. The companies in this report are extraordinary. The technology is extraordinary. But the same forces driving AI’s ascent — exponential compute, unchecked capital deployment, winner-take-most dynamics — are also generating a set of risks that most mainstream coverage either skirts around or buries in the small print. We want to detail out the real hurdles humanity faces as AI systems grow more capable, and get frank about what they are — so experts can start devoting more time to solving for them.

What follows is a unified picture: the physical limits on what AI can build, the human costs of what it displaces, and the policy and trade contests over who gets to shape it. These aren’t separate stories. They’re the same story — and understanding how they connect is the only honest way to assess where we’re headed.


1. THE SILICON CEILING

AI capability no longer scales with software innovation alone. It scales with access to chips. Training compute for frontier models has been doubling roughly every five to six months. The largest AI supercomputers now hold more than 200,000 chips — hardware that costs billions of dollars before a single model is trained.

The market is alarmingly concentrated. NVIDIA-designed GPUs form the backbone of almost all frontier training infrastructure globally. The H100 itself is now mostly a unit of measurement: since mid-2025, Blackwell chips have made up more than 90% of new NVIDIA compute,

and their successor, Vera Rubin, began reaching cloud providers in the summer of 2026. By Epoch AI’s estimates, cumulative compute capacity grew from roughly 2.4 million H100-equivalents in Q1 2024 to more than 20 million by

the end of 2025. NVIDIA accounted for about two-thirds of it, down from four-fifths at the start of 2024. Google’s own TPUs make up most of the rest, at close to a fifth; AMD, Amazon and Huawei together account for about 15%.



Source: Epoch AI, Data on AI Chip Sales (August 2026)

In response, the hyperscalers are building their own silicon. Google’s TPUs, AWS’s Trainium, Microsoft’s Maia — these are production infrastructure, not research prototypes. The driver isn’t incremental efficiency. It’s strategic vulnerability. Depending on a single dominant supplier for the most critical input in your industry is a risk no serious operator should accept indefinitely.

In the Middle East, Abu Dhabi’s Mubadala is the majority owner of GlobalFoundries, with about 73% after a May 2026 share sale. GlobalFoundries is a US-headquartered contract chipmaker with fabs in the United States, Germany and Singapore and 2025 revenue of US \$6.8 billion, but it makes mature-node chips, not the leading-edge silicon that AI accelerators need. Saudi Arabia has built out its National Semiconductor Hub (NSH), with the aim to attract at least 50 fabless semiconductor design companies to the Kingdom by 2030 and

a deep-tech venture fund of more than SR1 billion (about US \$266 million) to provide early-stage funding. As of mid-2025, 20 companies were already in various stages of setting up.

Building a full semiconductor fab from scratch — the kind that produces cutting-edge AI chips at 3nm or 5nm — takes US \$20 billion or more and several years even for the industry leader: TSMC’s first Arizona fab took about four and a half years from announcement to volume production. It also takes a talent ecosystem that doesn’t yet exist at scale in the region. Saudi Arabia has the capital, energy, and scale to build compute-adjacent manufacturing and packaging, but the consensus from industry analysts is that the region is most realistically targeting chip design, packaging, and assembly in the near term — not frontier fabrication.

For now, the Gulf’s most direct path to advanced silicon runs through

Washington. In November 2025, the US Commerce Department authorised the export of up to 35,000 Blackwell GB300 chips each to HUMAIN and G42 — establishing the legal and regulatory mechanism for further approvals. That’s a meaningful unlock, but a gigawatt of AI capacity takes 400,000 to 500,000 GPUs by NVIDIA’s own reckoning: Jensen Huang has put the 10 gigawatts NVIDIA is building with OpenAI at four to five million. Against that, 35,000 chips is well under a tenth of a gigawatt, which illustrates how much of the region’s AI ambition still depends on external silicon supply.

The consequence is structural. The firms that can secure large-scale accelerator capacity can train larger systems, iterate faster, and sustain performance advantages. AI capability is increasingly a function of hardware access — and hardware access is increasingly a function of geoeconomic alignment.

SPOTLIGHT | THE CHALLENGES AHEAD IN THE GLOBAL AI RACE

2. THE ENERGY PARADOX

The efficiency case for AI in industrial systems is real. Google DeepMind’s deployment of machine learning across its data centre cooling infrastructure cut cooling energy by 40% — a 15% reduction in overall energy overhead at an already highly optimised facility. Across heavy industry, AI-driven process optimisation has delivered 5–15% reductions in energy use. Predictive maintenance cuts unplanned downtime by 20–30%. In electricity grids, AI forecasting improves wind and solar output prediction accuracy by 10–30%, giving operators the confidence to run higher renewable penetration without sacrificing stability.

And then there’s the paradox.

Global data centre electricity consumption rose 17% in 2025, to about 485 TWh, and demand from AI-focused data centres grew by roughly half. The IEA expects total data-centre demand to roughly double again to 950 TWh by 2030, with AI-focused sites tripling. The strain is already visible in legacy

“A single AI-focused hyperscale data centre consumes as much electricity as 100,000 households. The largest ones under construction are expected to use 20 times that.”

IEA, 2025

markets. In Ireland, data centres used 23% of metered electricity in 2025, up from 21% in 2023, according to the Central Statistics Office; earlier forecasts of 32% by 2026 now look too high. In Virginia — home to the world’s densest concentration of data centres — they absorbed 26% of the state’s entire electricity supply in 2023.

The deeper problem in these markets isn’t demand. It’s infrastructure age. In the US, 70% of transmission lines and power transformers are more than 25 years old, according to the Department of Energy. The median new power project now waits more than five years from grid-connection request

to operation (Berkeley Lab), and US data-centre developers, the IEA notes, are turning to on-site gas generation rather than wait. In the PJM market — spanning Illinois to North Carolina — capacity prices for 2025/26 jumped more than ninefold on the prior year, and all three auctions since have cleared at PJM’s price cap, most recently in July 2026, with data centre growth cited as the primary driver. Households in Ohio and Maryland are already seeing their electricity bills rise as a direct result. Europe faces its own version: planning constraints, fragmented grids, and competing decarbonisation timelines are slowing the pace at which new AI infrastructure can come online.

The Middle East is a structurally different story. The Gulf doesn’t inherit these constraints — it builds around them. Energy costs in Saudi Arabia run at roughly half the UAE rate, and Qatar’s lower still, making the region among the most cost-competitive environments globally for power-intensive compute. Qatar, the UAE, and Saudi Arabia are some of the world’s largest energy providers — making them natural locations for large-scale data centres not just because of what they produce, but because of what they’re building now.

Regional data centre capacity is set to triple from 1GW in 2025 to 3.3GW by 2030, according to PwC. The national AI campuses announced since run to several gigawatts more on paper; financing and connection to power, not ambition, will set the real pace. Critically, that infrastructure is being built from scratch to AI-era specifications — liquid cooling, high-density rack configurations, and AI-optimised power distribution designed in from day one, not retrofitted. Dubai’s Moro Hub, certified by Guinness World Records as the world’s largest solar-powered data centre, shows how clean power can be designed in as well. Saudi Arabia’s HUMAIN is targeting 1.9 gigawatts of data centre capacity by 2030, and the 5-gigawatt UAE–US AI campus announced for Abu Dhabi in May 2025 would be the largest outside the United States.


Geography reinforces the structural advantage. The Middle East sits at the intersection of Europe, Africa, and Asia — a natural routing point for latency-sensitive applications serving the Global South. Sovereign Wealth Funds shorten the financing cycles that delay private developers elsewhere, although the

region’s flagship AI platforms, G42 and HUMAIN, are now also courting outside investors. And governments are offering long-term power contracts, land grants, and special economic zones that sharply compress deployment timelines.

None of this means the position is frictionless. Clean energy, including nuclear, made up close to 28% of the UAE’s energy mix at the government’s last count in 2024, against a target since raised to 35% by 2030–31. Cooling in a desert climate remains a genuine engineering challenge. And chip access runs through channels that no sovereign wealth fund can fully control. A fourth consideration becomes more important as capacity expands: operational resilience. AI infrastructure

needs not only sufficient power, but a dependable supply, backed by reliable cooling, connectivity and recovery systems. Building that resilience into new facilities adds to upfront costs, but is essential to keeping services running as demand grows.

Economists call this the Jevons Paradox: the more efficient a technology becomes, the cheaper it gets to use — so people use more of it, and total consumption rises. AI is a textbook case. But how painful that dynamic becomes depends on the infrastructure you’re starting with. If your grid was built in the 1960s, it’s a crisis. If you’re building from scratch in 2026, it’s an engineering challenge — and a solvable one, provided resilience is designed in from day one.



Data-centre electricity demand to 2030. China and the US are forecast to account for nearly 80% of global growth. Source: IEA, Energy and AI (2025), via Nature



SPOTLIGHT | THE CHALLENGES AHEAD IN THE GLOBAL AI RACE

3. WHO CONTROLS THE COMPUTE TIER

Hardware and energy are necessary but not sufficient. Sustaining frontier AI requires the infrastructure that binds them: hyperscale data centres, high-capacity networking, long-term energy contracts, and the capital to hold it together under conditions of extreme operational complexity. These requirements are narrowing the field.

Between 2021 and 2026, combined capital expenditure by Amazon, Microsoft, Alphabet, and Meta is set to rise from roughly US \$130 billion to more than US \$720 billion a year on the companies’ own guidance: growth of about 40% a year for five years, and nearly five times what the four spent as recently as 2023. Part of the latest increase is price rather than capacity: both Microsoft and Amazon have cited the rising cost of memory chips in lifting their 2026 plans. The United States holds roughly three-quarters of global AI computing power. China holds about 15%.



TOP 4 HYPERSCALER CAPITAL EXPENDITURE, 2021–2026
USD BILLIONS, CALENDAR YEARS

Company	2021	2022	2023	2024	2025	CAGR 2021-25	2026 Guidance
amazon	61	64	53	83	132	21%	~220
Microsoft	27	28	41	76	118	44%	~175*
Alphabet	25	31	32	53	91	39%	195-205
Meta	19	32	28	39	72	39%	130-145
Total	132	155	154	251	413	33%	720-745 +74-80% vs 2025


■ Growth on reported actuals ■ Company guidance, July 2026 (not included in CAGR)

* Microsoft guided to about \$190 billion until a July 2026 lease-accounting change moved some data-centre leases out of reported capex; it said its underlying plans were unchanged.
Note: Microsoft converted from its June fiscal year to calendar years; Microsoft and Meta include finance leases. CAGR = compound annual growth rate, 2021–2025. Totals may not add due to rounding.
Source: Company 10-K and 10-Q filings (2021–2025); company guidance with Q2 2026 results. July 2026. Lucidity Insights analysis.

Company filings and analyst estimates

The chart below shows how computational share has evolved by country since 2019 — and how durable American dominance has proved, despite sustained Chinese investment. Private companies now own about 80% of the world’s AI supercomputing capacity, up from 40% in 2019, by Epoch AI’s count. This is a structural departure from every previous era of computing, when academic and government labs played a central role in setting both the pace and the direction of progress.

SHARE OF AGGREGATE AI COMPUTATIONAL PERFORMANCE (16-BIT FLOP/S) BY COUNTRY, 2019–MAY 2025
US dominance is durable: China is the only sustained challenger.



As infrastructure scales, advantages compound. Larger operators iterate faster, absorb energy market volatility, negotiate better supply contracts, and deploy new capacity more rapidly. This isn’t a temporary head start. It is a moat that deepens with each investment cycle — and the barriers it creates for smaller firms, research labs, and emerging markets are not easily overcome.

Governments are responding accordingly, treating compute infrastructure the way previous generations treated ports and power lines: as foundational national capability. The EU has committed over €7 billion through EuroHPC, which now oversees 19 AI Factories and in July 2026 opened bidding for up to seven “AI Gigafactories” expected to draw more than €20 billion of private investment. India’s IndiaAI Mission allocates over US \$1.2 billion toward domestic compute and has put more than 38,000 GPUs at the disposal of Indian start-ups and researchers at subsidised rates. In the Gulf, sovereign-backed AI facilities are embedded directly within national economic transformation strategies — not as innovation policy, but as long-horizon infrastructure bets.

SPOTLIGHT | THE CHALLENGES AHEAD IN THE GLOBAL AI RACE

4. THE EMPLOYMENT DISRUPTION

In May 2025, Anthropic CEO Dario Amodei issued an unusually blunt warning: AI could eliminate half of all entry-level white-collar jobs within one to five years, spiking unemployment to between 10% and 20%, the top of that range a level not seen since the Great Depression.

The data so far supports a narrower version of that warning. US employers announced more job cuts in the fourth quarter of 2025 than in any fourth quarter since 2008, according to outplacement firm Challenger, Gray & Christmas, and attributed 54,836 cuts that year to AI. In the first eight months of 2026, AI-attributed cuts more than doubled to 116,175, or 22% of all announced cuts, and AI was the most-cited reason for five straight months. Stanford’s Digital Economy Lab finds no widespread, economy-wide displacement, but employment among 22- to 25-year-olds in the most AI-exposed occupations now sits 19% below where it would be had it kept pace with less-exposed jobs. Amazon cut 14,000 corporate roles in October 2025 and another 16,000 in January 2026, the largest workforce reduction in its history; AI featured in how the company framed the cuts, though CEO Andy Jassy has insisted they were “not even really AI-driven — not right now, at least.” The World Economic Forum projects 92 million jobs displaced globally by 2030, with 170 million new roles created

— a net gain that exists only at the macroeconomic level and tells individual displaced workers nothing useful.


What makes this wave different from every previous wave of automation: it isn’t targeting physical labour. It’s targeting cognitive labour. Lawyers, analysts, marketers, coders, consultants. The bottom rung of almost every knowledge profession is being automated first. The career ladders that generations used to climb are being cut away from the ground up — and the people most affected are the ones who haven’t yet had time to climb them.

The diffusion gap compounds this. The UNCTAD data below is stark: across every dimension that matters for AI readiness — ICT infrastructure, R&D

capacity, industrial base, skills, and finance — the gap between developed economies and least-developed countries is not closing.

The countries most exposed to AI-driven displacement of their labour forces — where outsourcing and low-cost knowledge work are significant economic contributors — are precisely the countries least equipped to build domestic AI capacity or retrain workers into AI-adjacent roles. The productivity gains accrue to those who already have the infrastructure. The disruption lands hardest on those who don’t.

AI is not a neutral rising tide. It is a capability concentrator — and right now, the concentration is accelerating.



Frontier technology readiness by country grouping. LDCs trail across Skills, ICT, R&D, Industry, and Finance. The gap is structural, not cyclical. Source: UNCTAD, Technology and Innovation Report



“We, as the producers of this technology, have a duty and an obligation to be honest about what is coming. I don’t think this is on people’s radar.”

Dario Amodei
CEO, Anthropic

5. THE BIAS PROBLEM NOBODY WANTS TO NAME

AI systems learn from historical data. And historical data reflects historical prejudice. When Amazon trained a recruiting AI on its own hiring records — historically male-dominated — the system began downgrading résumés containing the word “women’s.” It wasn’t programmed to discriminate. It learned to. The COMPAS algorithm, used by US courts to assess recidivism risk, was found to wrongly flag Black defendants who did not go on to reoffend as high-risk at nearly twice

the rate of white defendants. A 2025 Stanford-led study, published in Nature, found that when ChatGPT generated and then scored more than 34,500 résumés, it gave older men the highest ratings, even when their résumés were built from the same information as women’s.

What makes this dangerous at scale is the aura of objectivity that algorithms carry. When a human rejects your application, you can challenge it. When an algorithm does, it feels mathematical.

Final. Discrimination baked into code operates across millions of simultaneous decisions — far more consequential than any individual human bias, and far harder to surface, audit, or reverse.

This isn’t a niche problem. It’s the default condition of systems trained on unrepresentative data and deployed without adequate oversight. And it’s compounding quietly, at scale, right now.

6. THE AI RACE

For most of the past three years, the global AI race has been written as a two-player game. But this isn’t a bilateral race any longer. The Gulf — and more specifically, the UAE and Saudi Arabia, are increasingly playing a major role in the development and financing of AI.

Start with the numbers. In 2026, Amazon, Alphabet, Microsoft and Meta alone are guiding to more than \$700 billion in capital expenditure, most of it for AI — more than the entire Apollo program in today’s dollars. The core battleground is compute: the US holds 74% of global AI processing power; China holds 14%. Washington has spent nearly four years managing that gap through chip export controls. China’s answer was efficiency: DeepSeek, a frontier model built under severe hardware constraints, wiped \$600 billion off NVIDIA’s market cap in a single day in January 2025. Constraints drive innovation in ways trade policy doesn’t always anticipate.

While the US and China have competed on chips and compute, the Gulf has moved with unusual clarity —

establishing itself as a third centre of AI capital and capacity.

In November 2025, the US authorised up to 35,000 Blackwell GB300 chips each for Saudi Arabia’s HUMAIN and Abu Dhabi’s G42. The US President’s May 2025 visit to the Gulf had already reframed the region’s position: less a customer of American technology, more a commercial partner in deploying it. The UAE’s MGX is targeting \$100 billion in AI infrastructure. HUMAIN is pursuing a \$77 billion compute buildout. These are not passive financial bets. They are deliberate moves to acquire relevance in a world where compute capacity has become geo-economic currency.

The logic is straightforward: the US wants trusted international partners to host compute at scale. The Gulf needs the technology and partnerships to execute its economic transformation at speed. That alignment is proving durable — and it’s reshaping who gets a seat at the table when the rules of the AI era get written.


Meanwhile, China is building its own stack. In early 2025, DeepSeek’s own testing reportedly put Huawei’s Ascend 910C at roughly 60% of an NVIDIA H100’s performance on inference, though well behind on training, and Huawei launched its next-generation Ascend 950PR in March 2026. Beijing has ordered state-funded data centres to use only domestic AI chips, and to strip foreign chips out of projects less than 30% complete. The split between the two technology supply chains is becoming structural.



SPOTLIGHT | THE CHALLENGES AHEAD IN THE GLOBAL AI RACE



TIMELINE OF AI REGULATIONS AROUND THE WORLD



Beyond the US and China, more than 80 countries now have formal AI strategies, according to the World Bank. Three governance architectures are competing for dominance. The EU AI Act, which became law in August 2024, introduces binding risk-tiered obligations with material penalties. Its bans and its rules for general-purpose models already apply, but in July 2026 the EU pushed its core high-risk obligations back to December 2027 and

August 2028. The US framework pivots around competitiveness and speed. China integrates AI oversight within existing cybersecurity and content regulation. The UK, through its Safety Summit and AI Security Institute, has tried to carve a middle lane — only to join the US in refusing to sign the International AI Declaration at the Paris AI Action Summit in February 2025.

Neither bloc is likely to emerge as a clear winner, and the real risk is a race to the bottom, where speed trumps safety everywhere at once. The policy timeline above shows just how quickly governance has gone from voluntary guidance to a geoeconomic instrument.

Source: Lucidity Insights

SPOTLIGHT | THE CHALLENGES AHEAD IN THE GLOBAL AI RACE

7. THE QUESTIONS WE'RE NOT READY TO ANSWER

The ethics of AI aren't abstract. Who owns the data AI is trained on? Who is liable when an AI makes a wrong medical diagnosis? What happens when AI-generated content makes truth indistinguishable from fabrication? The EU AI Act is the most comprehensive framework to date. South Korea's AI Basic Act took effect in January 2026. But in the US — home to the world's most powerful AI companies — the Trump administration has actively resisted oversight in favour of competitive speed, and in December 2025 ordered federal agencies to challenge state AI laws in the name of a single national framework. The result is a global race to capability with no agreed rules of the road.

And then there's the question that sits behind all the others.

Geoffrey Hinton, Nobel laureate and the figure most responsible for the deep learning revolution now reshaping civilisation, estimates a 10–20% probability that AI leads to human extinction within 30 years. In February 2026, Amodei put the odds of a “country of geniuses in a data center”



arriving within ten years at 90%, and called a much faster arrival “more like a 50/50 thing.” Jensen Huang, who in 2024 put AGI about five years away, declared in September 2026 that “AGI has arrived” with OpenAI's GPT-6 Astra; critics said the claim came with no evidence and no definition. A survey of 2,778 AI researchers, conducted in late 2023, found 38–51% gave at least a 10% chance of extinction-level consequences. There is no consensus. Yann LeCun, who left Meta at the end of 2025 to found his own lab, argues today's language models are not a path to human-level intelligence at all. 76% of the 475 researchers surveyed for the AAI's 2025 presidential panel say scaling up current approaches is

unlikely to produce general intelligence. The sceptics are credible, and worth hearing.


The companies in this report are among the most powerful institutions in human history. The decisions they make over the next five years will shape not just the global economy, but the nature of work, knowledge, and power itself. Silicon constraints and energy bills and export controls and employment statistics are all real and immediate. But they are also proximate causes sitting in front of a more fundamental question: what kind of future are we actually trying to build, and who gets to decide?



“We don't need to believe the worst case is probable to believe it deserves serious attention. We buy insurance for risks far less likely.”

Sam Altman
CEO, OpenAI

ETHICAL AI GOVERNANCE ARCHITECTURE



Source: Lucidity Insights Research & Analysis

THE BOTTOM LINE

The physical constraints on AI are not temporary friction. They are the conditions within which the next decade of development will actually unfold. Compute bottlenecks that scale with hardware intensity. Energy infrastructure that cannot keep pace with digital ambition. Regulatory architectures pulling in incompatible directions. Employment displacement concentrated in the people and places least equipped to absorb it. An AI trade war where the rules are being written and rewritten in real time.

For founders and investors reading

this in the Gulf, the context matters in a specific way. This region sits at an interesting intersection: capital-rich, energy-abundant, with newer infrastructure and explicit sovereign intent to build domestic AI capacity at scale. The window is real. But it sits inside a global system that is becoming less open, less predictable, and less evenly distributed by the month.

AI optimises what exists. It cannot substitute for the judgment, the governance, and the accountability structures that don't yet exist. Building those — urgently, honestly, and with

a longer time horizon than the next funding round — is the challenge that matters most.



DUBAI'S AI POSITIONING

Dubai has positioned itself as a global bridge in the AI economy, embedding AI at the highest policy level with the world's first Minister of State for Artificial Intelligence in 2017. Through initiatives like Dubai AI Campus at DIFC, the city is translating strategy into applied infrastructure — connecting capital, regulation, and enterprise deployment. Rather than competing solely on frontier models, Dubai is winning on orchestration and speed. The Dubai AI Festival, has become a critical global convening point for the AI ecosystem.

DUBAI'S NEXT BET: AGENTIC AI AS THE NEW GROWTH ENGINE

Dubai Isn't Waiting for the AI Future — It's Already Building It.

In the final days of April 2026, Sheikh Hamdan bin Mohammed bin Rashid Al Maktoum, Crown Prince of Dubai, began issuing a series of directives with far-reaching implications: Dubai's government — and its entire private sector — was to transition to Agentic AI within two years.

The directive was not framed as another exploratory programme or limited pilot. It was presented as an economy-wide shift, extending the emirate's AI ambitions from government departments into the companies that power its economy.

That decision captures something important about Dubai's approach to artificial intelligence. Plenty of governments have published AI strategies; far fewer are attempting to reorganise the machinery of government and commerce around the technology. Dubai is increasingly placing itself in the second camp.

From Vision to Operating Model

The foundations were laid well before generative AI became a permanent fixture in boardroom conversations. In 2017, the UAE became the first country to appoint a Minister of State for Artificial Intelligence, naming His Excellency Omar Sultan Al Olama to the role. The appointment was widely regarded at the



time as symbolic. In retrospect, it looks more like an early marker of intent. Two years later, the UAE Cabinet adopted the National Artificial Intelligence Strategy 2031, setting out its ambition to become a global AI leader by the end of the decade.

There is a substantial economic case behind that ambition. PwC estimates that AI could contribute close to 14% of UAE GDP by 2030, equivalent to roughly US\$96 billion. Dubai's own

modelling is similarly striking. The Dubai State of AI Report, published by Digital Dubai and the Dubai Future Foundation, estimates that AI implementation could generate more than AED235 billion, or approximately US\$64 billion, in cumulative economic impact for the emirate by 2030. On those projections, Dubai alone would account for roughly two-thirds of the country's total AI-driven economic contribution—an outsized share for one of seven emirates.

The projections matter, but the numbers only tell a part of the story. More revealing is the consistency with which Dubai has translated policy into institutions, infrastructure and implementation—and its apparent unwillingness to allow anything to slow that process.

From Smart Government to Intelligent Government

Dubai's AI push is the latest chapter in a digital transformation that has been underway for more than two decades. Government services first moved online and then onto mobile applications, while data was progressively connected across departments.

Artificial intelligence builds on that foundation, but the ambition is considerably larger. In April 2024, Sheikh Hamdan launched the Dubai Universal Blueprint for Artificial Intelligence, or DUB.AI. The plan was designed to accelerate adoption across government and the wider economy, attract AI companies and talent, and support the Dubai Economic Agenda D33 target of generating AED100 billion a year through digital transformation.

The first phase combined institutional leadership with ecosystem development. Dubai appointed Chief AI Officers across government entities and the Dubai International Financial Centre (DIFC) introduced the first initiative under this blueprint, the Dubai AI Campus. Billed as one of the largest dedicated clusters for AI and advanced technology companies in MENA, the campus has already attracted more than 505 companies to its facility - achieving its target two years ahead of schedule. This initiative has already created more than 2,020 jobs.

Inside government, adoption is already moving beyond generic productivity tools. In 2025, the inaugural Dubai State of AI Report assessed more than 100 high-impact use cases that were either operating or under development across urban planning, healthcare, finance, mobility, procurement,

recruitment, compliance and strategic decision-making. Applications range from predictive healthcare and mobility digital twins to automated business licensing, fraud detection and AI-powered procurement and auditing.

One initiative illustrates the practical value of this approach particularly well. Digital Dubai's Dubai Population Now platform combines AI, machine learning and predictive analytics to produce a continuously updated estimate of how many people are present in the emirate—not merely how many officially reside there. The system draws on connected government records covering areas such as residency, transport, utilities, healthcare, property, education, tourism and the public-sector workforce.

For a city as transient as Dubai, the use case is critical. The resident population reached 4.58 million at the end of 2025, yet the average number of people present during peak daytime hours was 6.39 million. That additional 1.81 million people — roughly 40% of the resident base — creates real demand for roads, hospitals, schools, housing and other services. A continuously updated

population model gives planners a view of that demand that a conventional census, refreshed only periodically (every few years), cannot provide.

This is where Dubai's AI programme becomes more consequential. The objective is not simply to place new tools in the hands of government employees; it is to change how government anticipates needs, allocates resources and delivers services. Hamad Obaid Al Mansoori, Director General of Digital Dubai, has described AI as “not a future concept” but an everyday force enabling “a proactive, personalised, and predictive government.” A digital government responds when a request is submitted. A predictive government identifies the need earlier. An agentic government could increasingly act on it, proactively.

The UAE's Next Bet: Agentic Government

In April 2026, the UAE raised the stakes by announcing a plan to transition 50% of government sectors, services and operations to agentic AI within two years. Unlike the chatbots and copilots that have defined the first wave of mainstream generative AI,



agentic systems can monitor conditions, interpret changes, recommend a course of action and execute a sequence of tasks with greater autonomy.

Sheikh Mohammed bin Rashid Al Maktoum, Vice President and Prime Minister of the UAE and Ruler of Dubai, described AI as the government’s “executive partner” in supporting decisions, improving services, increasing operational efficiency and refining results in real time. That language goes well beyond the familiar promise of workplace productivity. It positions AI as part of the operating architecture of government itself.

The accompanying framework calls for policies, processes and procedures to be redesigned around AI capabilities, with ministers and directors general assessed on the speed and effectiveness of adoption. Each federal entity was initially given 90 days to move a selected service through exploration, design and implementation planning. The 4 May directive then carried that ambition into Dubai’s private sector. Under Sheikh Mohammed’s direction, the Emirate set a two-year transition horizon, with specialised training through business councils, government-funded incubators for agentic AI companies, and helping establish dedicated investment vehicles. Two weeks later, the UAE Cabinet approved a federal governance framework defining responsibilities across ministries and government entities. It also launched what it described as the largest training programme in UAE Government history, covering 80,000 employees—from ministers and senior executives to new joiners—and approved an initial package of agentic AI applications spanning services for citizens, residents, businesses and the wider public.

By early June, the scale of the private-sector plan had become clearer: 295,000 Dubai companies were to be supported over two years, alongside the development of 100 specialised AI assistants and the establishment of 50 agentic AI companies.



Individually, each announcement is ambitious. Taken together, they reveal a coordinated architecture for adoption across federal government, Dubai’s private sector and the institutions responsible for governance, training and investment. Even more telling, is that the sequence unfolded over approximately six weeks. Dubai continued to set targets and assign responsibility for delivering them.

For AI founders, this creates a market worth watching and engaging with. Governments are among the largest purchasers of technology in any economy. When public institutions move from experimenting with AI to redesigning workflows around it, new categories of demand emerge. If businesses across an entire city are being encouraged to follow, the addressable opportunity expands further—from specialised agents and integration tools to cybersecurity, governance, training and sector-specific infrastructure.

Building the Infrastructure Around the Ambition

Mandates alone do not create an AI economy. Dubai has also been assembling the physical, commercial and human infrastructure needed

to support one, much of it before the latest agentic AI targets were announced.

The Dubai AI Campus at DIFC provides one part of that foundation. In addition to workspace and a dedicated commercial licence for AI-dependent businesses, it offers access to computing capacity, accelerator programmes and partnerships with major global technology companies. Its value lies not only in the number of firms it can host, but in bringing capital, talent, infrastructure and prospective customers into the same ecosystem. Talent development is the other half of the equation. The Dubai AI Academy, launched in April 2025 on the final day of the inaugural Dubai AI Week, aims to educate 10,000 emerging and established leaders. Its partners include Oxford University’s Saïd Business School, Udacity and the Minerva Project, while its AI for Civil Service programme makes training available to government employees. Rather than a standalone education initiative, the academy is intended to create a pipeline of people capable of implementing AI across institutions and industries.

Dubai AI Week 2025 provided a global showcase for that agenda, attracting more than 30,000 participants from

over 100 countries across 10 major events. More than 30 initiatives, partnerships and agreements were announced between government entities, investors, startups and private companies. The Global Prompt Engineering Championship drew nearly 3,800 participants competing for AED1 million in prizes, while students from 60 Dubai schools took part in AI Week in Schools.

The 2026 edition takes place on 26–27 October at the Dubai World Trade Centre, with the aim of doubling the inaugural event’s audience and attendees.

Regulation Is Becoming Part of the Infrastructure

The faster AI moves into healthcare, finance, government services and public infrastructure, the more urgent questions of security, transparency and accountability become. Dubai’s response has been to treat regulation not simply as a constraint on the industry, but as another component of the ecosystem it is trying to build.

In January 2025, the emirate launched the Dubai AI Seal, a certification framework for trusted AI companies. By May, 325 companies operating in Dubai and representing 77 international offices had applied. The programme uses a six-tier grading structure, with its highest designation recognising businesses making the strongest economic and technological contributions to Dubai’s AI sector. Early recipients included IBM and e&, one of the region’s largest telecommunications groups. The commercial significance of the certification increased in October 2025, when Dubai Finance directed government entities to award AI-related service contracts only to companies holding the Seal. For suppliers seeking public-sector work, a voluntary trust mark had effectively become an entry requirement.

Dubai’s AI Policy for Government Entities, introduced in April 2025, addresses the other side of the

market. Built around explainability, interoperability, human-centricity and proactive governance, it establishes a common framework for how public bodies should develop and deploy AI. Together, the policy and the Seal show an emirate attempting not only to attract more AI companies, but to define what a credible supplier and a responsible government customer should look like.

For startups, that raises both the cost of entry and the potential reward. Meeting Dubai’s evolving standards will require documentation, governance and technical safeguards. Companies that can satisfy those requirements, however, gain access to a government and business market being pushed towards adoption from the top.

What Does This Mean for Founders?

Every major technology cycle creates a period in which governments, companies and consumers are searching for new solutions at the same time. AI is firmly in that phase. What makes Dubai unusual is that its government is attempting to act simultaneously as customer, regulator, investor, talent developer and ecosystem builder—and to coordinate those roles at speed. None of this guarantees that the targets will be met. The harder questions now concern execution. How many pilots will become scaled deployments? How easily will startups gain access to public procurement? Can governance keep pace with systems that are becoming more autonomous? And can Dubai attract and retain enough technical talent to meet the demand its policies are creating?

Those questions remain unresolved, as they do in every market pursuing AI at scale. What Dubai has demonstrated more clearly is a willingness to move from ambition to action. Some of the most consequential commitments discussed in this report reflect a long-term approach to developing the infrastructure, talent and partnerships needed to advance AI at scale. The pattern is familiar: Dubai is investing

in the capabilities that will shape its future.

For founders, investors and policymakers watching from Singapore to São Paulo, the more interesting question may not be whether Dubai delivers every target exactly as promised. It may be what the Emirate’s willingness to make its largest bets early, and at speed, says about the kind of AI market it intends to become.

As Sheikh Hamdan has put it, the objective is straightforward: to make Dubai the world’s leading city for the economic and commercial adoption of advanced AI. The work now is to turn that objective into an operating reality.



UAE RANKS 34TH IN GLOBAL INDUSTRY 5.0 INDEX — BUT PLACES TOP 10 FOR AI-DRIVEN RESILIENCE



A new global benchmark suggests the next phase of economic competition will not be defined by technology alone — but by how well countries use technology to strengthen society, sustainability, and resilience.

According to the newly released Industry 5.0 Index, developed by the Oliver Wyman Forum in partnership with the University of California, Berkeley, the United Arab Emirates ranks 34th globally among 92 countries preparing for what many economists are calling the Fifth Industrial Revolution. But the headline number tells only part of the story. The UAE ranks 9th globally for resilience, placing it firmly among the world's most technologically prepared economies when it comes to infrastructure, cybersecurity, and the ability to withstand economic shocks. In the September 2026 Global Financial Centres Index update, Dubai was ranked first globally for FinTech, and it ranks seventh overall, the highest of any financial centre in the Middle East, Africa and South Asia.

THE RISE OF INDUSTRY 5.0

Industry 5.0 represents the next evolution of the digital economy.

Where Industry 4.0 focused on automation, AI, and productivity, the Industry 5.0 framework shifts the focus toward human-centric growth — using emerging technologies not just to improve efficiency, but to enhance wellbeing,

environmental sustainability, and economic resilience. The index evaluates countries across 30 indicators grouped into three core pillars:



Source: Oliver Wyman Forum & UC Berkeley, Industry 5.0 Index

Countries that perform well across all three dimensions are expected to lead the next phase of economic transformation. According to the report, success in this new model could add as much as \$1 trillion annually to global GDP, beyond the productivity gains already expected from emerging technologies.

UAE: A TECHNOLOGY-FIRST GROWTH STRATEGY

The UAE's strong showing in the resilience pillar reflects more than a decade of aggressive national investment in digital infrastructure and artificial intelligence.

Few countries have moved as quickly to position themselves as global testbeds for emerging technologies.

The government's National AI Strategy aims to embed artificial intelligence across the economy, with projections that AI could generate \$91 billion in economic value and contribute up to 20% of the country's non-oil GDP by 2031.

At the same time, Abu Dhabi has announced plans to become the world's first fully AI-native government by 2027, automating large portions of public administration and transitioning services to sovereign cloud infrastructure.

Infrastructure is another major differentiator.

The UAE is investing heavily in large-scale AI data centers, next-generation telecom networks, and national cybersecurity frameworks, including plans for a nationwide 6G network by 2030.

Projects such as the planned Stargate AI data-center initiative are designed to serve not just the domestic market, but users across a 3,000-kilometer radius, positioning the country as a regional digital hub.

AI ADOPTION IS ALREADY HAPPENING

One of the report's most striking findings relates to workforce adoption.

According to Oliver Wyman Forum research, 83% of UAE residents report using AI tools at work at least once per week, significantly higher than the 61% global average.

This suggests AI adoption in the UAE is not confined to tech specialists or large corporations — it is already permeating the broader workforce.

Meanwhile, government programs such as the National AI

Summer Camp and new university initiatives are accelerating the development of an AI-ready talent pipeline.

THE NEXT CHALLENGE: TALENT AND SUSTAINABILITY

Despite its strong infrastructure and digital capabilities, the UAE's overall ranking of 34th highlights areas for continued improvement, particularly in workforce development and sustainability indicators.

Industry 5.0 leaders tend to combine cutting-edge technology with deep human capital and green innovation — areas where advanced European economies currently dominate the rankings.

Still, the trajectory is clear. The UAE has positioned itself as one of the most ambitious adopters of AI and digital infrastructure in the world, and the Industry 5.0 Index suggests that strategy is already translating into global competitiveness.

In the emerging Industry 5.0 era, the countries that win will not simply deploy the most technology. They will be the ones that use it to build smarter, greener, and more resilient economies.



Q&A

Mohammad Alblooshi
CEO of DIFC Innovation Hub

What distinguishes Dubai's approach to AI as the world moves from experimentation to execution?

The question facing governments, businesses and investors is no longer whether to adopt AI, but how quickly they can turn adoption into measurable value. What sets the UAE apart is that it has made AI part of how the nation operates. This began in 2017, when the UAE appointed the world's first Minister of State for Artificial Intelligence. That commitment has continued under the leadership of His Highness Sheikh Mohammed bin Rashid Al Maktoum, Vice President and Prime Minister of the UAE and Ruler of Dubai, and the directives of His Highness Sheikh Hamdan bin Mohammed bin Rashid Al Maktoum, Crown Prince of Dubai, Deputy Prime Minister and Minister of Defence of the UAE. This year, the UAE committed to transitioning 50 per cent of government services and operations to agentic AI within two years. Dubai has extended that ambition to the private sector, with plans to support 295,000 companies to adopt AI. The foundations are already in place: more than 70 per cent of working-age adults in the UAE use generative AI, and the GCC reports the highest enterprise AI adoption of any region, at 84 per cent.

What role are DIFC and Dubai AI Campus playing in translating AI adoption into economic value?

DIFC is at the heart of this shift. As we advance our ambition to become the world's first AI-native financial centre, we are embedding AI across

our regulatory frameworks, operations, talent and infrastructure. That transformation is expected to generate USD 3.5bn (AED 12.9bn) in economic value and create 25,000 jobs. Dubai AI Campus brings founders, investors, enterprises and regulators together within one of the world's most trusted financial jurisdictions. It surpassed its target of 500 AI companies two years ahead of schedule and is now home to 505 AI companies, which have created more than 2,020 jobs. For global AI firms, Dubai is not simply a market. It is a launchpad to scale across the Middle East, Africa and South Asia, and on to the world.

What needs to happen next for organisations to scale AI responsibly and effectively?

Globally, nine in ten organisations now use AI, yet fewer than half have scaled it. The advantage will belong to those that close this gap first. The next phase will be defined less by the capability of the models themselves than by the institutions that put them to work. That means regulation that builds trust, capital that funds scale and talent that turns intelligence into outcomes. Financial centres sit at the intersection of all three. Those that act now, by embedding intelligence into how markets are governed, financed and run, will set the standards that others follow. The journey from adoption to execution will demand decisive leadership, meaningful investment and trusted governance. That is why platforms such as Dubai AI Festival matter: they turn dialogue into the partnerships that take AI forward.

SPOTLIGHT | DUBAI AI CAMPUS

DUBAI AI CAMPUS:
THE INSTITUTIONAL BET
TO MAKE DUBAI A GLOBAL
AI LAUNCHPAD



Artificial intelligence is no longer constrained by model capability. It is constrained by where it can scale. Dubai’s answer to that constraint is the Dubai AI Campus — a deliberate attempt to institutionalize AI inside a live economic system.

Launched within the Dubai International Financial Centre (DIFC) in May 2024, the Campus was designed not as another innovation cluster, but as applied AI infrastructure. Its ambition is clear: attract over 500 AI companies by 2028 (which it achieved 2 years ahead of schedule), and mobilize hundreds of millions in AI-related capital into and through Dubai.

The Campus aims to position Dubai as a global centre for applied AI, with a focus on translating innovation into practical deployment across the economy. Globally, most AI hubs are either

research-heavy or venture-driven. Few integrate regulation, capital, enterprise demand, and cross-border scalability in one environment. Dubai is attempting to compress that stack into a single platform.

Why DIFC Changes the Equation

Location is not incidental. At the end of H1 2026, DIFC was home to 10,018 active registered companies, up 30% year on year, operating under an independent English common law framework trusted by global investors. In H1 2026 its financial ecosystem included 1,134 regulated financial services firms, 327 banks and capital markets firms, and 592 wealth and asset managers, including the region’s highest concentration of hedge funds. Within that base, the number of AI, FinTech and innovation firms reached 1,933 in H1 2026, up 39 per cent year on year, and DIFC Innovation Hub

welcomed 361 new companies in the first half of the year. For AI companies, that means immediate proximity to capital and enterprise clients — not theoretical access. “DIFC provides an institutional advantage, not just a location,” Alblooshi notes.

In a world where AI scale increasingly depends on regulatory clarity and capital intensity, operating inside a globally recognized legal and financial jurisdiction reduces friction in fundraising, structuring, and international expansion. This is infrastructure — not branding.

A Hybrid Ecosystem by Design

The Campus was intentionally built as a hybrid. Today it serves global AI firms entering the Middle East, UAE-based startups scaling internationally, and enterprises deploying AI solutions across regulated industries.

Global firms contribute technology and intellectual property, local startups bring agility and regional knowledge, and enterprises provide demand and practical use cases. Together, they create a foundation for commercial adoption. Early traction has been strongest in FinTech, climate and sustainability solutions, PropTech, enterprise productivity AI, and smart city and government applications. These sectors align closely with Dubai’s economic strengths in finance, infrastructure, logistics, and public-private collaboration.



كامبس دبي للذكاء الاصطناعي
DUBAI AI CAMPUS



The Campus is not trying to outbuild Silicon Valley on frontier models. It is embedding AI into sectors that already move capital and policy at scale.

Embedded Capital

One of the most underappreciated aspects of Dubai AI Campus is how deeply capital is integrated.

Investors — venture capital, corporate venture arms, sovereign and institutional funds — are not invited in occasionally. They are structurally embedded in the ecosystem through daily proximity, structured programs, and major platforms such as Dubai AI Festival.

“Investors are embedded into the ecosystem rather than engaged episodically.”

That reduces one of the biggest bottlenecks in AI scaling: timing. In many ecosystems, companies build first and seek capital later. In Dubai, the objective is simultaneity — company formation, capital access, regulatory clarity, and enterprise partnerships happening in parallel. The long-term ambition is not just company count. It is economic output. DIFC’s planned transformation into the world’s first AI-native financial centre is expected to generate USD 3.5 billion (AED 12.9 billion) in economic value and an anticipated 25,000 jobs.

Regulation as a Growth Lever

As global AI regulation tightens — from Europe’s AI Act to emerging governance frameworks in the U.S. and Asia — companies increasingly need compliant environments that still move quickly. Dubai’s model is regulated acceleration. The Campus supports regulatory experimentation, policy dialogue, and real-world implementation feedback as governance frameworks evolve. This positioning matters. In 2026, the constraint on AI is no longer model intelligence. It is institutional integration: compliance, liability, cross-border data governance, and enterprise readiness. Dubai is betting that clarity beats chaos.

Dubai as a Launchpad — Not Just a Market

Dubai’s role in the global AI economy extends beyond its appeal as a regional market. For AI companies seeking international growth, the city offers a base from which to enter new markets, build cross-border partnerships, and scale globally.

That role builds on Dubai’s long-established position as a bridge between East and West. For decades, the city has connected Asia, Europe, and Africa through trade, logistics, capital flows, and financial services. AI is becoming the next layer of this international network.

For companies operating in regulated sectors or requiring cross-border expansion into emerging markets, Dubai offers speed of setup, investor confidence, and enterprise access within a single jurisdiction. That orchestration is the point.

The Real Strategic Play

Dubai AI Campus is not competing on raw model training scale. It is not trying to dominate chip manufacturing or frontier lab research. Instead, it is positioning Dubai as a global scaling node in the AI economy — a place where companies formalize structure, raise capital, deploy into enterprise, and expand into high-growth markets.

If successful, by 2028 the Campus will not merely host 500 companies. It will embed AI into finance, climate systems, infrastructure, and governance — turning applied intelligence into economic export. In a world where frontier capability is consolidating among a handful of global players, smaller jurisdictions cannot win on compute. But they can win on orchestration.

Dubai’s bet is that orchestration — of capital, regulation, and enterprise demand — is the next competitive frontier. If that thesis proves right, Dubai AI Campus will be less a hub and more a hinge point in the global AI economy.



SPOTLIGHT | DUBAI AI FESTIVAL 2026



DUBAI AI FESTIVAL 2026:
A CONVERGENCE OF
CAPITAL, GOVERNANCE
AND DEPLOYMENT



This Special Report is the official publication of the Dubai AI Festival, a DIFC Initiative, which returns in October 2026 at the Dubai World Trade Centre for its third edition. Held under the directives of His Highness Sheikh Hamdan bin Mohammed bin Rashid Al Maktoum, Crown Prince of Dubai and Deputy Prime Minister of the UAE, the Festival is organised by Dubai AI Campus in partnership with DIFC and the Office of the UAE Minister of State for Artificial Intelligence, Digital Economy and Remote Work Applications. Since its launch in 2024 as the Dubai AI & Web3 Festival, the platform has expanded steadily in both scale and institutional weight. The inaugural edition drew more than 6,800 delegates from over 100 countries; the second, in April 2025, drew more than 10,000 participants, including 500 investors. The move to DWTC for 2026 reflects this upward trajectory, with organisers targeting more than 20,000 attendees, 200+ global speakers and 200+ exhibitors.

20,000+

Attendees

500+

Investors

200+

Global Speakers

200+

Exhibitors

120+

Countries Represented

150+

Focused Sessions

The Festival is positioned as a business-focused convening where governments, enterprises, investors and technology leaders meet in a single environment. Its structure combines public-stage dialogue with curated networking and commercial engagement formats, allowing participants to move from discussion to direct interaction.

Agenda & Thematic Focus

The programme spans two days and is structured across Plenary and Innovation stages, alongside roundtables and workshops. Topics range from national AI strategies and global governance to industry transformation and applied enterprise use cases. The Festival is organised around eight themes: Infrastructure & Investment; Global AI Governance; Innovation & Ethics; AI Literacy; Future of Work;

AI for All; AI Cornerstone; and Immersive Experiences.

Day One centres on global leadership and economic positioning, including discussions on AI as a national strategy, regulatory coordination and large-scale partnerships. Day Two shifts toward sectoral deployment and operational integration, with sessions examining AI in banking, life sciences, sustainability, autonomous systems and customer experience.

A dedicated governance track, AI Nations, forms a distinct component of the 2026 programme. Hosted in collaboration with Dubai Future Foundation, AI Nations convenes ministers, regulators and policy leaders to explore responsible AI frameworks, interoperability standards and cross-border cooperation. The inclusion of this platform places governance dialogue alongside commercial activity rather than separating it into a parallel forum.

CONVERGING GLOBAL FUTURETECH LEADERS

20+ GLOBAL UNICORNS

Revolutionising industries with cutting edge AI innovations

20+ GOVERNMENT AGENCIES/MINISTRIES

Coming together to drive AI adoption & innovation in governance

ENTERPRISES


Presenting transformative AI solutions

20+ TOP VENTURE CAPITALISTS

Gathering to explore investment opportunities to advance AI development

Source: Dubai AI Festival / Dubai AI Campus

DUBAI AI FESTIVAL: GROWTH ACROSS EDITIONS (2024-2026)



CONCLUSION: THE INFRASTRUCTURE DECADE OF INTELLIGENCE

“If 2023 was the year of AI awakening, 2024 the year of acceleration, and 2025 the year of AI Agents, then 2026 is the year of reckoning.”

Artificial intelligence is no longer a frontier experiment. It is no longer a feature inside productivity software. It is becoming infrastructure — embedded into capital markets, healthcare systems, supply chains, national security architectures, and the operating models of firms large and small. The question has shifted from “Can it work?” to “Who can scale it — and who will be left behind?”

The data throughout this report tell a consistent story. Adoption is widespread: nearly nine in ten organizations now use AI regularly in at least one business function. Copilots remain the operational baseline. Agentic systems are beginning to execute, not just assist. Frontier models continue to advance at breakneck speed. And yet, meaningful enterprise transformation remains concentrated in a minority of firms. Only 44% of organizations say AI is scaling across the enterprise; only 37% see any EBIT impact from it.

The bottleneck is no longer intelligence. It is integration. Infrastructure, energy, data quality, regulatory clarity, and

talent depth are the gating variables of this decade. AI compute demand is rising faster than many grids can support. Semiconductor supply chains remain sensitive to trade policies. Governance regimes are diverging — from the EU’s risk-tiered AI Act to the United States’ speed-first posture and China’s state-integrated oversight. The Middle East is charting a fourth path: state-coordinated acceleration without heavy-handed constraint, embedding AI into national transformation agendas while tying its compute build-out to American chips and cloud platforms. This is no longer just a technology race; it is a systems race — and increasingly, a sovereignty race.

Frontier model development remains concentrated in a handful of countries, with the United States and China operating at scale. But participation in the AI economy is broadening. Open-weight models, led by Chinese labs, now trail the closed frontier by months, not years. Sovereign strategies are emerging. The Middle East, particularly the UAE and Saudi Arabia, has positioned itself not as a

passive adopter but as an orchestrator — investing in compute, cloud regions, regulatory frameworks, and capital platforms designed to attract and scale AI companies globally.

Money flows confirm the shift. Capital is moving up the stack — toward chips, data centers, and energy — signaling that AI is being financed as economic infrastructure, not short-cycle software. Amazon, Alphabet, Microsoft and Meta alone are guiding to more than \$700 billion in capital expenditure this year. Anthropic’s valuation more than doubled in under four months, to \$965 billion; OpenAI’s nearly tripled in a year, to \$852 billion. And some of that capital now moves in a circle, as NVIDIA helps finance the data centers that buy its chips. The companies dominating this era are not simply building applications; they are shaping the physical and institutional foundations of intelligence.

But the gains are arriving unevenly. Early signs of divergence are visible — between large enterprises and SMEs, between advanced and emerging economies, between firms redesigning

workflows and those layering AI onto legacy systems. At companies with more than \$1 billion in revenue, 40% now say they are scaling AI agents, up from 27% a year earlier; at smaller firms, the figure is stuck at 22%. Without deliberate institutional redesign, AI risks amplifying existing inequalities.

At the same time, the ethical questions have become more urgent. Governance has shifted from voluntary principles to enforceable law, though the EU has since pushed its rules for high-risk AI back to December 2027 at the earliest. The dual-use nature of AI — capable of accelerating drug discovery and cybercrime alike — demands sophistication, not fear, and certainly not complacency. This September alone, a drug developed with generative AI entered a Phase III trial, and OpenAI placed its newest model in the top tier of its own cyber-risk framework.

Meanwhile, the physical constraints on AI — compute, energy, chips, grid capacity and, increasingly, operational resilience — are not temporary

friction. They are the conditions within which the next decade will unfold. The decisions made now about infrastructure investment, regulatory design, and talent development will shape who controls intelligence at scale.

That is the reckoning: capital, capability and risk have arrived at once, and most institutions are not yet organized to absorb them. The State of AI in 2026 is clear: we have entered the infrastructure decade of intelligence. AI is reorganizing the global economy in real time. The institutions that align capital, compute, governance, and human capability will compound advantage. Those that hesitate will not simply fall behind technologically — they will fall behind structurally.

The future is not being built in laboratories alone. It is being engineered in boardrooms, ministries, data centers, and policy chambers. And, as we all already know — either intuitively or through real-world experience — this AI race is already well underway.




Lucidity Insights does two things: it reports on the region’s most important technology and investment stories with the depth they deserve, and it builds the tools that help founders act on them. Our special reports inform the ecosystem. Our Founder Hub — the region’s first all-in-one platform for startup founders — equips it: with investor matchmaking, market intelligence, and the fundraising playbooks founders need to build, raise, and scale across the Middle East and Africa.

This Report

This is the official publication of Dubai AI Festival 2026, which was produced in collaboration with Entrepreneur Middle East, Dubai AI Campus, and DIFC, and was published in October 2026.

Our Reports & Publications:



Don't Miss a Report
Subscribe to our free weekly newsletter – and get the latest reports, analysis, and insights from across the regional startup ecosystem, straight to your inbox.
lucidityinsights.com/reports

Commission a Special Report

Interested in sponsoring or commissioning data-driven special report for your sector, market or audience? We partner with leading organizations across the region. Contact us today:
partnerships@lucidityinsights.com

Editorial Team

Chief Content Officer | Erika Masako Welch
Editorial | Ashok Raman, Pierrick Ribes

Published In Partnership With



All rights reserved 2026. Opinions expressed are those of the contributors. All research, data analysis, and reporting for this special report was conducted by Lucidity Insights. Entrepreneur Middle East and all subsidiary publications in the MENA region are officially licensed exclusively to BNC Publishing by Entrepreneur Media Inc. No part of this publication may be reproduced without the written permission of the publisher.

Follow Us

- lucidity-insights
- @LucidityInsight
- @lucidityinsights
- @lucidityinsights
- @lucidityinsights



MEA IPT Unicorn Tracker

Track the region's billion-dollar startups



lucidityinsights.com/unicorn-tracker

**THE STATE OF AI IN 2026:
HOW THE MIDDLE EAST IS BUILDING THE FUTURE OF INTELLIGENCE**