INDUSTRY VISION REPORT 2026

AI: The New Essential Infrastructure

$700B	$80B	$5.6B	77%	50%	80.8%
capex 2026	IaaS spend	mkt 2026	AI deploy	lock-in risk	GenAI growth
AI INFRA	SOVEREIGN AI	QUANTUM	6G	CLOUD	ENT. AI

$480B	$80B	28.7%	2029
Global AI Spend 2026	Sovereign Cloud IaaS	Quantum Market CAGR	First Commercial 6G
UBS Research	*Gartner, Feb 2026*	*Research & Markets*	*3GPP / IDC*

→ Inference > Training for first time in 2026

→ Sovereign cloud growing 35.6% YoY

→ 70%+ of Fortune 500 running quantum pilots

→ 6G architecture locked in at MWC 2026

CONTENTS

EXECUTIVE SUMMARY

Five Infrastructures Defining the Next Decade of Enterprise Technology

The strategic technology decisions made in 2026 will shape enterprise competitiveness for the rest of the decade. Across AI infrastructure, quantum computing, next-generation networks, cloud architecture, and enterprise AI adoption, a single theme has emerged: infrastructure is no longer a commodity. It is a strategic asset.

$480B	$80B	$5.6B	50%
Global AI spend 2026	Sovereign cloud IaaS 2026	Quantum market size 2026	Enterprises citing lock-in as critical
UBS, 2026	*Gartner, Feb 2026*	*R&M, 2026*	*SUSE, Mar 2026*

Five conclusions stand out from our analysis:

→ AI infrastructure has crossed the threshold from discretionary investment to national strategic priority, with sovereign AI spend growing 35.6% year-on-year.

→ Quantum computing is transitioning from academic prototyping to enterprise experimentation — over 70% of Fortune 500 companies now run quantum pilot programmes.

→ 6G architecture decisions made today will determine who controls the AI inference layer of tomorrow's networks — this is a sovereignty question, not a technology question.

→ Cloud vendor lock-in has become the defining infrastructure risk — 50% of global enterprises cite it as a critical or major challenge.

→ Enterprise AI has entered its production phase — the question is no longer whether to adopt, but how to govern, scale, and extract measurable value.

CHAPTER 1

AI Infrastructure & Sovereign AI

The hyperscaler capex arms race has reached a scale that rivals the GDP of mid-sized economies. A structural shift is underway — from centralised, vendor-controlled AI compute towards sovereign, locally governed infrastructure.

$700B	57.5%	35.6%	$140B+
Hyperscaler AI infra spend 2026	On-premise share of AI infra spend	YoY growth in sovereign cloud IaaS	Middle East sovereign AI pledges
Tech-Insider, 2026	*Mordor Intel., 2026*	*Gartner, Feb 2026*	*Mordor Intel., 2026*

The Scale of the Commitment

Global AI infrastructure spending approaches $480 billion in 2026, with hyperscalers committing nearly $700 billion in combined capital expenditure. Amazon leads with a projected $200 billion in total capex — a 50% increase from 2025. Google and Alphabet announced $175-185 billion, nearly doubling their 2025 commitment. Microsoft Azure revenue grew 39% year-on-year in constant currency.

On-premise infrastructure accounts for 57.5% of AI infrastructure spending, driven by data-residency mandates and regulatory frameworks including HIPAA and NIS2. The majority of enterprise AI infrastructure continues to sit outside hyperscaler clouds.

Sovereign AI: From Strategy to Procurement

Worldwide sovereign cloud IaaS spending is forecast to reach $80 billion in 2026 — a 35.6% increase year-on-year. Middle East and Africa (89%), Mature Asia-Pacific (87%), and Europe (83%) are recording the highest growth rates. Saudi Arabia and the UAE alone have pledged over $140 billion to build domestic AI hyperscale campuses.

> *"The scaling of AI inside enterprises is becoming a major driver of sovereignty. As organisations become more ambitious, many will find that their existing infrastructure architectures cannot provide the assurances required for sensitive or regulated workloads."*
>
> — Kate Hanaghan, Chief Research Officer, TechMarketView

The NVIDIA Moat and Its Challengers

NVIDIA captured roughly 80% of 2025 AI accelerator revenue and maintains a 4-million-developer CUDA ecosystem. Hyperscalers are responding with custom ASICs — Google TPU v6e, AWS Trainium2, Microsoft Maia 100 — projected to account for 20% of training hours by 2026, compressing NVIDIA pricing power by up to 30% for bulk orders.

Strategic Implication

Enterprises building AI infrastructure in 2026 face an irreversible architecture choice: hyperscaler dependency with scale advantages, or sovereign infrastructure with control, compliance, and long-term strategic optionality. The decision locks in for at least five years.

CHAPTER 2

Quantum Computing Readiness

Quantum computing is at an inflection point. The technology must now demonstrate scalable commercial use cases to justify the billions already invested — and the window for early-mover advantage in enterprise quantum strategy is narrowing.

$5.6B	28.7%	>70%	$90B
Quantum market size 2026	Market CAGR 2026-2032	Fortune 500 running quantum pilots	Quantum market by 2035
R&M, 2026	*R&M, 2026*	*Ent. QC Report, 2026*	*McKinsey, 2024*

From Research to Enterprise Experimentation

The global quantum computing market expanded from $4.39 billion in 2025 to $5.59 billion in 2026. Over 70% of Fortune 500 companies are running quantum pilot programmes, with 160+ enterprise projects globally. The dominant use cases are optimisation (40%) and machine learning experimentation (30%). Cloud-based quantum access accounts for nearly 80% of enterprise usage.

Hardware: The Race to Fault Tolerance

IBM released its industrial-scale quantum blueprint in June 2026, targeting the world's first large-scale fault-tolerant quantum computer by late 2029. Google's Willow chip demonstrated scalable error correction with approximately 105 qubits. Microsoft Azure Quantum assembled a multi-modal ecosystem combining Quantinuum, IonQ, and Atom Computing. Superconducting qubits dominate with 41% market share.

> *“When institutions like Aalto University own their quantum computers, it means their data, their IP, and their expertise stay theirs. That's not a feature — that's a strategic posture.”*
>
> — Jan Goetz, CEO & Co-founder, IQM Quantum Computers

National Quantum Strategies

Japan announced a $7.4 billion quantum commitment in 2026. The EU Quantum Flagship allocates EUR 1 billion over ten years. The US Department of Energy committed $2.5 billion. Finland, ranked the number two global quantum cluster by ECIPE, installed AaltoQ20 — a 20-qubit IQM-built system giving university students direct hands-on access to real quantum hardware.

Strategic Implication

Enterprises that begin quantum workforce development and use-case mapping now will hold a meaningful advantage when commercially relevant quantum advantage arrives before 2030. The window to build internal expertise is approximately five years.

CHAPTER 3

6G & Network Sovereignty

6G is no longer a distant promise. The architectural decisions being standardised today will determine who controls the intelligence layer of the world's networks for the next two decades.

2029-30	77%	38%	$14B
Expected first 6G commercial launches	Operators expecting faster 6G AI deploy	CAGR of AI-in-telecom through 2030	AT&T Open RAN programme value
3GPP / IDC, 2026	*Nvidia / MWC 2026*	*Covalense, 2026*	*Covalense, 2026*

MWC 2026: The Architecture Decision

MWC Barcelona 2026 was the moment AI-native networking moved from aspiration to evidence. NVIDIA secured commitments from over a dozen global operators — including BT Group, Deutsche Telekom, Ericsson, Nokia, SK Telecom, SoftBank, and T-Mobile — to build 6G on open, AI-native, software-defined platforms. AT&T;'s $14 billion Open RAN programme targets 70% of traffic on open platforms by 2026.

The Sovereignty Dimension

6G architecture is a sovereignty battleground. The choice of 6G vendor in 2026-2027 will determine who can observe, analyse, and potentially control the AI inference layer embedded in future network infrastructure. 6G technical specifications now include quantum-resistant security as a design requirement.

> *"6G is not merely an iteration of mobile technology. It is the infrastructure that will distribute AI across devices, the edge and the cloud."*
>
> — Borje Ekholm, President & CEO, Ericsson — MWC Barcelona 2026

Strategic Implication

Enterprises that treat 6G network infrastructure as a commodity procurement decision are making a sovereignty and security decision by default. The vendors running your 6G network will have structural visibility into your AI workloads.

CHAPTER 4

Cloud Platform Lock-In

Cloud vendor lock-in has graduated from a procurement concern to a board-level risk. As AI workloads deepen dependency on proprietary hyperscaler ecosystems, the cost and complexity of switching is rising faster than most enterprises anticipated.

50%	$1T	20%	39%
Enterprises citing lock-in as critical or major	Global cloud market value 2026	Workloads shifting to local via geopatriation	US enterprises concerned about lock-in
SUSE Survey, Mar 2026	*Holori, 2025*	*Gartner, 2026*	*SUSE Survey, 2026*

The Lock-In Ratchet

The global cloud market exceeds $1 trillion in 2026. AWS holds 30% global market share, Azure 20%, Google Cloud 13%. As enterprises embed AI workloads into hyperscaler ecosystems — training on proprietary chips, deploying inference through proprietary APIs — the practical cost of switching escalates dramatically.

A March 2026 SUSE survey of 596 global enterprise technology leaders found 50% cite vendor lock-in as a critical or major challenge. US enterprises are the most concerned globally at 39% — outpacing the 25% global average. 59% are moving to hybrid cloud and 16% to private cloud in response.

The Sovereign and Neocloud Response

Gartner estimates 20% of current workloads may shift to local providers through geopatriation. OVHcloud surpassed EUR 1 billion in annual revenue in FY 2025. Neocloud providers — Nscale, CoreWeave, Carbon3ai — are positioned specifically for sovereign AI and HPC workloads. Forrester predicts 2026 is the year governments adopt tech nationalism in AI procurement.

> **Strategic Implication**
>
> The window to architect for cloud portability is closing. Enterprises that have not mapped critical workloads to open-standard infrastructure will find the exit cost rising in direct proportion to their AI investment.

CHAPTER 5

Enterprise AI Adoption

Enterprise AI has entered its production era. The dominant challenge is no longer access to AI capability — it is governance, integration, and the extraction of measurable business value from AI investments already made.

$6.15T	80.8%	40%	24%
Global IT spending in 2026	GenAI model spend growth 2026	Enterprise apps with AI agents by end 2026	Enterprises citing AI as critical challenge
Gartner, Feb 2026	*Gartner, Feb 2026*	*Market Clarity, 2025*	*SUSE, 2026*

From Pilot to Production

Global IT spending reaches $6.15 trillion in 2026. Data center spending jumps 31.7% to top $650 billion. Generative AI model spending grows at 80.8%. AI cloud infrastructure reaches $37.5 billion, with inference spending ($20.6 billion) exceeding training ($16.9 billion) for the first time — enterprises have moved beyond experimentation into production deployment at scale.

The Governance Gap

Despite the pace of adoption, a governance gap persists. SUSE found that implementing AI ranked as a critical (24%) or major (37%) challenge globally. 97% of telcos recognise gaps in AI talent, and only 42% offer skills training. The enterprises extracting maximum value are those with production-grade AI governance and board-level accountability.

Agentic AI: The Next Wave

By end of 2026, 40% of enterprise applications are forecast to include AI agents. Orange has announced an agentic AI framework entirely LLM-independent, ensuring sovereignty over its models. The shift to agentic AI is compressing the timeline between AI capability and business impact — but governance frameworks for autonomous AI systems are lagging the technology.

> *“AI investment is shifting toward business-critical domains — revenue generation, product innovation, and competitive differentiation — not isolated efficiency experiments.”*
>
> — IDC Japan AI Infrastructure Report, March 2026

Strategic Implication

The enterprises extracting maximum value from AI in 2026 have deployed production-grade AI infrastructure with clear measurement, ownership, and accountability at the board level — not pilot governance frameworks.

LOOKING AHEAD

Ten Key Signals for 2026

01 **Inference overtakes training**
AI cloud spend on inference exceeds training for the first time — signalling the shift from AI development to AI deployment at enterprise scale.

02 **Sovereign cloud becomes standard**
$80B in sovereign cloud IaaS spend in 2026, growing 35.6% YoY. Europe will surpass North America in sovereign cloud spending by 2027.

03 **Quantum enters the enterprise pilot phase**
Over 70% of Fortune 500 companies are running quantum pilots. IBM targets fault-tolerant quantum by 2029. Talent pipeline is the binding constraint.

04 **6G architecture locks in at MWC 2026**
The NVIDIA-led AI-native 6G coalition secured 12+ operator commitments. The open vs. proprietary architecture battle will largely settle by 2027.

05 **Cloud lock-in reaches board level**
50% of enterprises cite vendor lock-in as critical. Geopatriation is real — 20% of workloads may shift to local providers within three years.

06 **Neocloud providers gain enterprise share**
CoreWeave, Nscale, and Carbon3ai are emerging as credible alternatives to hyperscalers for sovereign and HPC AI workloads.

07 **Agentic AI moves to production**
40% of enterprise apps will include AI agents by end 2026. Governance frameworks for autonomous AI systems are lagging the technology.

08 **Energy becomes the binding constraint**
Data center power demand is beginning to constrain AI infrastructure expansion. Liquid cooling accelerates as rack density surpasses 100kW.

09 **Geopolitics fragments the AI stack**
US-China chip export controls, EU AI Act implementation, and tech nationalism in government procurement are fragmenting the global AI landscape.

10 **Quantum-safe cryptography becomes urgent**
As quantum computing approaches commercial relevance, the window to migrate to post-quantum cryptographic standards is narrowing. 6G now requires it.

ABOUT STRAITHEAD

About Straithead

Straithead is an independent intelligence platform covering enterprise AI, quantum computing, cloud platforms, cybersecurity, networking, and digital infrastructure strategy.

We produce editorial analysis, research reports, and intelligence briefings for technology leaders, enterprise decision-makers, and the vendors who serve them. We accept no advertising, carry no vendor affiliations, and publish no sponsored editorial. Our analysis is independent by design.

Coverage: Enterprise AI — Quantum Computing — Cloud Platforms — Security & Networking — Analytics & Digital Transformation

straithead.com — hello@straithead.com

IMPORTANT DISCLAIMERS

Financial Disclaimer

Straithead is not a financial adviser, investment adviser, broker, dealer, or financial analyst of any kind. Nothing contained in this report constitutes financial advice, investment advice, investment research, a prospectus, or a recommendation to buy, sell, hold, or otherwise transact in any security, financial instrument, asset class, fund, or investment product.

All market size figures, revenue forecasts, growth projections, and financial estimates cited in this report are sourced exclusively from named third-party research organisations and are reproduced solely for informational and contextual purposes. **Readers must not rely on any information in this report to make financial, investment, or commercial decisions.** Any such decisions should be made only after consulting a suitably qualified and regulated financial adviser. Past performance is not indicative of future results. Market forecasts are inherently uncertain and subject to material change without notice.

Editorial Independence

This report is produced under Straithead's full editorial independence policy. No organisation, individual, vendor, or commercial entity has paid for inclusion in, or exercised any influence over, the content, conclusions, or recommendations of this report. All analysis and strategic conclusions are those of Straithead's editorial team alone.

Data Sources & Accuracy

Straithead has made reasonable efforts to verify the accuracy of all data cited. However, Straithead cannot guarantee the accuracy, completeness, or currency of any third-party research cited herein. All third-party forecasts and estimates are subject to change. Straithead accepts no liability for any errors or inaccuracies in third-party data reproduced in this report.

Copyright & Reproduction