



AI DECRYPTED:

A Guide for Navigating AI Developments in 2026



Stay ahead with us.



Introduction

In 2026, the central question is no longer whether models will improve, but who will control the infrastructure, energy, data, talent, and regulatory pathways that determine where artificial intelligence can be deployed at scale. As enterprises, hyperscalers, and sovereign actors race to lock in these advantages, the AI market is entering a phase where strategic positioning will matter as much as technical performance. **Ten key trends and challenges will influence AI in 2026.**

2025 marked a decisive inflection point as model developers and hyperscalers moved from experimentation to large scale consumer and enterprise deployment. New reasoning models were bundled into products for coding, research, and early agentic workflows, driving a surge in adoption of chatbots, personal assistants, and deep research tools. Awareness of embodied intelligence and humanoid robotics also grew, setting the stage for broader commercialization in 2026. User stickiness across platforms became a central battleground, especially in the U.S. and China.

Meta's late 2025 acquisition of Manus reflected how aggressively U.S. platforms are moving to secure agentic capabilities and paying users, while China's regulatory scrutiny of the deal highlighted a growing constraint on cross-border transfers of AI talent, models, and know-how. At the same time, the technology frontier continued to advance, with steady gains in coding, math, science, and long-horizon task execution. Deep reasoning models expanded use cases into research, analysis, and business operations, including generating presentations, tables, and complex graphics.


Rising demand for advanced compute drove hyperscalers, model developers, and hardware firms into massive energy, supply chain, and financing partnerships, accelerating data center buildouts while raising concerns about sustainability and circular financing. In parallel, China's open weight ecosystem surged after the DeepSeek Effect. DeepSeek's V3.2 release in December and the expected V4 in early 2026 reinforces China's efficiency-driven, diffusion-focused model, sharpening contrasts with the U.S. emphasis on capital-intensive frontier systems.

U.S.-China rivalry remained the dominant geopolitical backdrop. Partial rollbacks of export controls in response to China's rare earth restrictions set up intensified competition in third country markets in 2026, while power shortages and slow permitting in the U.S. pushed policymakers toward offshore compute and a more aggressive AI export strategy, including expanded deployments in the Middle East under the U.S. AI Action Plan.

As businesses aim to navigate the considerably changed AI landscape in 2026, we have outlined 10 emerging trends and challenges that are expected to define the future of AI this year.

Summary


1. **AI Bubble Fears Overblown:** Agentic AI, World Models, Embodied AI, and Enterprise Uptake Stoke Compute Demand
2. **U.S. AI Infrastructure Buildout Runs into Wall:** Lack of Energy, Other Key Inputs Stall Progress
3. **Trump Administration AI Initiatives Stretch Bureaucracy:** Finding Where to Put Effort Proves Challenging
4. **Regional AI Development Takes Off:** Middle East, Japan, India Emerge as Major Hubs, Along with Concept of AI Sovereignty
5. **Global Governance Fully Fractures:** U.S.-China Struggle to Lead Competing Blocs
6. **Open vs Closed Model Clash Escalates:** U.S. Labs Will Open More, but Chinese Labs Will Lead Diffusion
7. **China's Challenge Gets Complicated:** Alternative AI Stack Not Fully Ready
8. **Brussels Effect Hits AI Wall:** Europe Aspires to be Major Player
9. **U.S. Regulatory Efforts Hit Quagmire:** Federal and State Tensions Create More Confusion
10. **Advanced AI Diffusion Raises Risks:** AI-Driven Incident More Likely



AI Bubble Fears Overblown: Agentic AI, World Models, Embodied AI, and Enterprise Uptake Stoke Compute Demand

AI agents and multi-agent platforms will move from internal demos into daily enterprise and consumer use. Alongside wider deployment of multimodal models, early world models, and embodied AI through humanoid robots and autonomous vehicles, this will drive another surge in demand for advanced compute. Tokens will increasingly shift from mostly text to frames, views, and 3D representations, raising compute intensity for both training and inference while forcing the industry to build real scaffolding for safe, interoperable, and governable systems.

Compute drives electricity demand



Note: Hyperscalers will be adding power both behind the meter, building their own capacity, and from grid: Input shortages will impact both efforts
Source: IEA, EPRI, S&P Global, Goldman Sachs

Trend: Agents, world models, and science-based applications will drive a durable compute boom

AI in 2026 will look less like a speculative bubble and more like a sustained compute boom driven by agents, multimodal models, longer research loops, and rising inference demand beyond chatbots. In enterprises, agents will triage email, draft memos, pull background from news and filings, generate code, and hand off structured outputs to humans, with fastest adoption in software, image and video generation, and parts of healthcare and finance. Early world models from Fei-Fei Li's World Labs and Yann LeCun's new venture will begin shaping robotics, simulation, and scientific discovery.

Embodied AI will scale as advances in actuators, electronic skin, batteries, and spatially aware models push humanoid robots and autonomous vehicles into factory floors and logistics. On the consumer side, general purpose agents will be packaged for mass market use. Manus is an early example, combining planning, tool use, and task execution into a single interface, and Meta's acquisition shows that for major platforms, control of the agent layer is now as strategically important as control of the underlying model. In China, multi-agent frameworks are being embedded directly into smartphone operating systems, with ByteDance's ZTE prototype released in late 2025 showing how full app level control could work once regulatory and ecosystem constraints are resolved.

Leading models will approach or exceed top percentile coding in multi-file debugging, dependency resolution, and end to end integration, enabling early virtual engineering teams by late 2026.

Token volumes and demand for graphics processing units (GPUs), memory, and inference hardware will continue to rise, but efficiency gains will shift the main bottleneck toward power and energy rather than silicon.

Challenge: Interoperability and scaffolding lag behind capability

Multi-agent systems are advancing faster than the infrastructure for safe deployment, compliance, and cross-ecosystem operation. The Agentic AI Foundation launched in December by OpenAI, Anthropic, Microsoft, AWS, and Google to advance standards such as Model Context Protocol, Agents.md, and Goose is the first serious attempt to close this gap, but operating system (OS) level deployments like ByteDance's ZTE prototype show how far real-world capability is outrunning shared frameworks. Enterprises will need new layers for access control, observability, auditability, and incident response before agents can touch production systems, and embodied AI faces similar safety and liability gaps. Without shared protocols for context, tools, and identity, large-scale multi-agent systems will remain brittle and hard to govern even as capabilities accelerate.

Challenge: Security, safety, and governance risks from always-on agents

Always-on agents will introduce new and sometimes opaque attack surfaces. Enterprise agents with persistent access can be hijacked through prompt injection, compromised tools, or insider misuse. In high priority research and industrial domains such as energy, aerospace, advanced materials, biotechnology, and climate engineering, agent-driven acceleration may outpace existing safety,

export control, and classification frameworks. Diverging architectural paths will further complicate governance. Ecosystems focused on large, general-purpose models, and those emphasizing tightly orchestrated multi-agent stacks may not share compatible safety benchmarks or evaluation tools. Policymakers will almost certainly underestimate the operational burden of securing always-on, multi-agent infrastructures embedded deep inside corporate and government networks. Governments will begin to address these issues in earnest in 2026, with the U.S. Center for AI Standards and Innovation (CAISI) releasing a call for inputs around best practices for security and safety

around AI agents early this year and putting the issue high on its agenda.

Bottom line:


2026 will mark the shift of multi-agent platforms from pilots to everyday use, with world and multimodal models driving a sustained rise in global compute demand. Enterprises will begin handing real workflows to agents, leading models will reach top percentile coding on targeted tasks, and consumer devices will compete through OS-level assistants. The key question will be whether standards, governance, and security can mature fast enough to let these systems run inside critical workflows without creating unacceptable risk.

2

U.S. AI Infrastructure Buildout Runs into Wall: Lack of Energy, Other Key Inputs Stall Progress


Physical and supply chain limits on U.S. AI infrastructure will become unavoidable in 2026, pushing firms to look toward more energy rich regions to hedge against compute shortfalls. JPMorgan estimates U.S. companies signed about 9 gigawatts (GW) of AI data center capacity in 2025 on top of 4 GW in 2024, far short of the 100 to 300 GW projected by 2030. With no step change energy sources coming online in 2026, power, equipment, and permitting bottlenecks will increasingly slow new projects. AI power density has already outpaced grid readiness. A 1 GW campus hosting roughly one million GPUs often requires dedicated generation, water, and years of interconnection studies, while transformer and generator shortages further delay timelines. Projects like Texas' 1.2 GW Stargate site, which depends on new transmission infrastructure that may take years, show how grid and equipment constraints will push major capacity into 2027 or 2028 and drive up costs, creating opportunities for firms that can secure power or deploy interim solutions such as mobile generation and storage.

Global data-center electricity consumption, 2020–2030



Source: International Energy Agency (IEA), S&P

U.S. data-center grid power demand, 2020–2030



Trend: Gulf energy and capital emerge as the relief valve for U.S. AI compute

In 2026, the UAE, Saudi Arabia, and Qatar will become Washington's preferred partners for exporting U.S. aligned AI infrastructure as cheap energy, faster permitting, and sovereign capital absorb capacity the U.S. cannot bring online quickly. Led by Jacob Helberg and AI czar David Sacks, the Trump administration will accelerate approvals for advanced compute exports, allowing hyperscalers and their Gulf partners to scale from pilots to U.S. stack superclusters. The Pax Silica initiative, which Saudi Arabia, Qatar, and the UAE joined in early 2026, formalizes the Gulf's role inside U.S. aligned AI supply chains, signaling that offshore allied compute is now a core pillar of U.S. strategy.

A new AI Diffusion framework is expected to expand agreements with Abu Dhabi, Riyadh, and Doha, with safeguards on access, diversion, and monitoring. This will shift a growing share of U.S. linked training capacity offshore and create a long pipeline of power, cooling, and construction contracts tied to the U.S. AI stack, even as Washington debates how to balance domestic and offshore buildouts.

Challenge: Domestic politics and permitting slow U.S. buildouts

Federal efforts to accelerate approvals and clear equipment bottlenecks will face limits from state rules, local opposition, and rising public scrutiny as the year progresses and companies announce new deals and seek to host projects locally. As power bills rise and water use becomes more visible, protests and lawsuits targeting high-profile campuses will spread.

Developers will face higher community benefit requirements and longer timelines, leading some projects to be scaled back or quietly shelved. Intense market scrutiny here will increase pressure on project developers and investors.

Challenge: Offshoring compute creates new strategic dependencies




As late 2025 export licenses for GPUs and advanced networking gear to the Gulf begin to deliver capacity, pressure will grow in Washington to approve larger shipments. China-focused voices will push for tighter controls, while accelerationists Helberg and Sacks will argue that U.S.-aligned Gulf capacity is essential to maintaining a compute lead. By the end of 2026, major U.S. model developers are likely to rely on at least one Gulf campus for core training or redundancy, exposing them to future political shocks, export control disputes, or regional instability beyond U.S. jurisdiction.

Bottom line:

In 2026, power constraints and domestic politics will slow U.S. data center buildouts, while Gulf-based superclusters tied to the U.S. AI stack accelerate. The Trump administration will struggle to balance domestic infrastructure priorities with incentives to export the AI stack, an unfamiliar coordination challenge for the federal government. Investors and operators should treat U.S. campuses as longer-term, higher-risk projects and view Gulf capacity as both a near-term opportunity and a growing source of strategic dependence.

“This is benefitting the entire economy... We're seeing this huge infrastructure buildout, and it's benefitting people like plumbers, electricians, concrete pourers... it's really a broad-based boom.”

- David Sacks
White House AI Czar




3

Trump Administration AI Initiatives Stretch Bureaucracy: Finding Where to Put Effort Proves Challenging

In its first year, the Trump administration launched an unusually broad stack of AI initiatives, including the AI Action Plan, Project Genesis, Pax Silica, and related efforts such as the U.S. Tech Force. In 2026, these programs will move from strategy documents into early execution, testing whether the federal government can restructure itself, attract outside expertise, and manage initiatives spanning ramping up energy supply, critical minerals supply chain resilience, AI data center buildouts — including on federal land — government adoption of AI, and promotion of open-source and open-weight models. Some progress is likely, but flagship efforts will quickly collide with personnel, legal, bureaucratic, and political constraints.

U.S. federal AI implementation stack



Source: DGA-ASG Technology Policy Practice

Trends: From announcement to uneven execution

2026 will be the first true operational year of the AI Action Plan and its megaprojects, and progress will be uneven. Government controlled initiatives such as the Genesis Mission will move fastest as national lab data sets, simulation engines, and mission programs are consolidated into AI-driven research environments. Coordination will remain difficult as the White House Office of Science and Technology Policy, which authored the Action Plan, has limited staff while Departments of Commerce, State, and Energy control key authorities but pursue competing priorities. Pax Silica, launched just before the new year and led by Helberg at State, will compete with export controls, energy security, and diplomacy for senior attention.

Rapid private sector dealmaking and unclear ownership over infrastructure, permitting, and advanced compute exports will further complicate execution. Federal efforts to clarify the government's role in supporting industry, managing supply chain bottlenecks, deploying AI internally, and exporting the U.S. AI stack will make some progress in early 2026 but remain fragmented. Support for private buildouts will improve only at the margins, with modest easing of grid interconnection and permitting weighed down by legal challenges, bureaucratic inertia, and outdated regulations.

Federal agencies will deploy early mission-specific frontier models in enforcement, biosurveillance, and infrastructure monitoring, driving advances in materials science, climate modeling, and energy simulation.

The base case is visible but has uneven progress, with technical breakthroughs alongside growing friction around compute supply and interagency coordination.

Challenges: Physical, legal, and political constraints

Execution will prove far harder than planning. Power grid constraints, data center permitting delays, legal challenges, federal-local conflicts, and bureaucratic inertia will slow many high-profile government and private sector supported initiatives. Genesis will encounter delays as efforts to integrate sensitive scientific and environmental data sets run into classification rules, privacy protections, and agency-specific legal constraints. Resistance from departments reluctant to cede control over research pipelines will further slow consolidation. In practice, the pace and direction of AI progress in 2026 will be shaped as much by permitting offices, standards bodies, and regional politics as by technical performance.

Bottom line:

The AI Action Plan's expansive vision will meet operational reality in 2026 as overlapping initiatives are slowed by coordination failures and turf battles. Programs like the Genesis Mission are most likely to deliver early wins in materials, climate, and energy research, strengthening the case for centralized scientific compute. Whether these scale will depend more on the government's ability to clear grid bottlenecks, align agencies, and manage the politics of AI infrastructure rather than model advances.


4

Regional AI Development Takes Off: Middle East, Japan, India Emerge as Major Hubs, Along with Concept of AI Sovereignty

In 2026, AI competition will move from a U.S.-China duopoly toward a more multipolar landscape as regional hubs build their own stacks. While the U.S. and China will still dominate frontier research, more training, fine tuning, and deployment will shift to the Middle East, Europe, Japan, and India as governments pursue AI sovereignty. Countries are increasingly unwilling to rely solely on U.S. hyperscaler clouds for sensitive workloads, instead building domestic capability around open weight models and locally trained systems.

A new AI geography emerges

- AI Hubs - including model development/AI infrastructure
- Emerging AI hubs



Source: DGA-ASG Technology Policy Practice

Trend: Gulf superclusters become the main non-U.S. training belt for the U.S. AI stack

The Middle East is currently moving the fastest on compute infrastructure buildout. Following U.S. approvals for advanced GPU exports in late 2025, visible capacity will expand significantly throughout 2026. Abu Dhabi is on track to become one of the largest training hubs outside North America, while Saudi Arabia's industrial projects and Qatar's telecom-linked clusters provide cost-competitive campuses backed by cheap power and sovereign capital. For global firms, these hubs offer scalable alternatives for training and inference closer to clients in Europe and Africa, marking a transition in which major developers schedule core training runs in the Gulf rather than treating the region as a mere pilot site.

Trend: Japan and Europe position themselves as rules-aligned industrial anchors

Japan and Europe are positioning themselves as rules-aligned anchors in the global AI supply chain. Japan is strengthening its role as an infrastructure hub through production milestones in high-bandwidth memory and logic facilities, tightly integrating itself with allied chip supply chains. Its guidance-first regulatory approach aims to provide predictable rules for sensitive deployments in manufacturing and healthcare. Similarly, Europe is advancing "AI factories" and sovereign cloud offerings to provide regulated sectors with options that fall under EU data and safety laws. Major cloud providers are increasingly offering region-specific deployment options to meet these local sovereignty requirements.

Trend: India turns its scale into a regional AI services and deployment platform

India is leveraging its vast developer base and digital public infrastructure to become a regional AI platform for the Global South. The India AI Mission is pooling hardware resources into shared facilities, tying them to tenders that favor domestic cloud providers. Rather than focusing purely on frontier research, new Indian campuses are targeting the deployment of language models capable of handling large domestic language groups and public-sector tools tied to existing well-developed digital services platforms for identification and payments. While power and water constraints may limit the number of ultra-large campuses, India's primary role in 2026 will be as a global hub for services, fine-tuning, and the export of AI-driven administrative tools to other emerging economies.

Challenge: Fragmented, sovereignty-driven rules raise compliance and product complexity

As these regional hubs mature, multinational firms face a growing quagmire of conflicting mandates. Governments are increasingly attaching conditions to compute access, such as data localization requirements or specific model evaluation rules. Gulf partners often focus on security assurances and content controls, while Europe enforces strict liability regimes. Providers seeking to serve all markets simultaneously will need to maintain separate deployment configurations and distinct model versions to comply with these divergent sovereignty packages.

Furthermore, heavy reliance on regional hubs ties long-term AI strategies to local energy and water policies. Shifts in renewable buildouts or local resource stress could alter the cost equations of these data centers overnight. In the Gulf, regional security shocks remain a persistent risk to critical training capacity, while, in Europe and Japan, high energy prices and local opposition to data center construction could slow the scale of onshore options.

Bottom line:

2026 marks the point at which regional AI hubs become operationally critical for servicing emerging AI regional hubs. Firms that successfully orchestrate workloads across these diverse facilities — while managing the friction of divergent rules and local energy risks — will gain a significant resilience edge over competitors tied to a single jurisdiction.

“The more a society uses AI, the more we can influence the development of it...”

- Peng Xiao
CEO, G42


Challenge: Energy, water, and political risk at regional hubs will shape long-term exposure

Heavy reliance on Gulf or Indian data centers will tie long-term AI strategies to domestic energy and water policy, as well as the stability of relations with Washington. Subsidy changes, shifts in renewable buildout, or local water stress could alter cost equations quickly. In the Gulf, regional security shocks or U.S. sanction debates could threaten access to critical training capacity. In Europe and Japan, high energy prices and local opposition may slow some AI factory and data center projects, limiting the scale of onshore options for regulated sectors.

Global Governance Fully Fractures: U.S.-China Struggle to Lead Competing Blocs

Neither the U.S. nor China will succeed in leading a coherent global coalition, and interoperability across governance systems will become more difficult. China is likely to formalize its approach at the February India Impact Summit in Delhi, through the official launch of the World AI Cooperation Organization (WAICO) — proposed in July 2025. For many governments in Africa, the Middle East, Central Asia, and Southeast Asia, WAICO will offer a development-first alternative to Western regimes associated with export controls, security screening, and heavy compliance burdens. Western coordination around AI governance will narrow rather than expand in 2026. The Bletchley Park and New Delhi frontier AI safety tracks will run aground and drift toward technical exchanges rather than broad governance. India's AI Impact Summit, likely to largely exclude significant participation from the Chinese government, safety community, and AI company leaders, will reflect an effort to position Delhi as a Global South leader, with unclear support from developing countries that may tilt towards Beijing.

EU AI Act: Implementation Timeline (2025-2026)



Source: DGA-ASG Technology Policy Practice

Trends: Regionalization and bloc-based governance

Global AI governance efforts will shift toward regionally anchored systems that prioritize intra-bloc compatibility over universal standards. Beijing will increasingly bundle governance initiatives, standards of development, and infrastructure financing for Belt and Road countries through WAICO. Western efforts will become more technocratic and limited in scope, while middle powers experiment with hybrid frameworks drawing selectively from U.S., EU, and Chinese models. The Bletchley Park process is likely to see efforts to establish a permanent venue in a neutral country such as Switzerland or Singapore, while AI safety institutes in the U.S., UK, Singapore, Japan, and other developed countries continue to collaborate on establishing best practices for testing frontier models and emerging agentic systems.

Challenges: Fragmentation and uneven protection

Divergent rules on data localization, incident reporting, evaluation, and oversight will increase compliance costs for firms operating across regions and raise the risk of accidental violations. Governments with limited regulatory capacity may lock into rigid governance frameworks without clear enforcement or revision pathways. This fragmentation will also complicate responses to cross-border AI incidents and limit the sharing of tools, best practices, and trusted evaluation methods.

Bottom line:

Global AI governance will not converge toward a “minimally viable global governance framework.”

Instead, it will split into competing tracks featuring regional and plurilateral blocs, with China advancing a development-linked governance model through WAICO and Western governments relying on a patchwork of EU, U.S., OECD, and G7 frameworks to push leading Western AI labs to adopt new best practices and adhere to emerging standards. This fragmentation creates real risks but also opens space for middle powers and private actors to act as standards brokers. The core question is whether this fractured system can sustain enough coordination to manage shared AI risks that do not respect national borders.

6

Open vs Closed Model Clash Escalates: U.S. Labs Will Open More, but Chinese Labs Will Lead Diffusion

Competition between open and closed AI models is expected to reach a structural turning point. Rather than prevailing in one approach, the landscape will likely solidify into a dual-track ecosystem defined by geopolitical and commercial interests. U.S. frontier labs, facing intense pressure to maintain safety and secure revenue, will likely double down on high-assurance, closed-system deployments for their most capable models. Conversely, Chinese firms and European players, such as Mistral, are positioned to become the primary sources of competitive open-weight models for the global market. American laboratories will likely experiment with controlled releases and corporate-friendly licensing to blunt the progress of international rivals. While these firms offer stronger open-source versions of mid-tier models, their flagship systems will remain behind proprietary walls, accessible only through managed application programming interfaces. This "managed openness" serves a dual purpose: it protects the immense capital investments required for training while allowing firms to offer secure, compliant services to enterprise clients in sensitive sectors such as defense and finance.

Open-weight model diffusion accelerates in the Global South (2026)

Region	Why it adopts	Typical deployments	Binding constraints	What wins
Africa	Price + localization	Gov services; telecom; SME tools	Power + connectivity	Smaller open weights + edge-friendly
Middle East	Capital + power	Sovereign clouds; national platforms	Export-policy exposure	U.S.-aligned stacks + large clusters
South Asia	Scale + services	Outsourcing; agents in back office	Grid + reliability	Cost/performance + talent
Southeast Asia	Fragmented markets	Consumer apps; enterprise copilots	Data rules diversity	Local partners + tuned open models
Central Asia	Connectivity corridors	Gov + security use cases	Limited data centers	Host OF open models + regional hubs
Latin America	Enterprise modernization	Customer ops; fraud; public services	Skills + infrastructure	Managed deployments + open weights

Source: CSIS, DGA-ASG Technology Policy Practice

Trend: China as the gravity center for open-source model diffusion

Chinese firms will continue to move in the opposite direction, by positioning China as the global leader in advanced open-source weight models. DeepSeek's late 2025 V3.2 release, and accompanying research papers highlighting new hardware optimizations and its forthcoming V4 model aimed at coding performance, demonstrate how Chinese firms are pushing open weight systems directly into the most commercially valuable enterprise use cases. This accelerates the competitive pressure on U.S. closed model providers, especially in price-sensitive and emerging markets. Microsoft highlighted the spread of Chinese open weight models in the Global South in an early 2026 report, highlighting growing concern within the U.S. AI ecosystem about competition from Chinese models in developing markets. By promoting open weights through dedicated vehicles like WAICO, Beijing hopes to boost uptake of Chinese models in Belt and Road countries that will be receptive to a Chinese AI stack featuring Huawei hardware and capable open-source models to build at a significant cost differential from Western alternatives. This strategy will be particularly effective in the Global South, where U.S. models may be restricted by export controls or prohibitive pricing. In regions across Africa, Asia, and the Middle East, Chinese open-weight models are likely to become the default choice for public sector applications in education, healthcare, and agriculture.

Challenge: Fragmented ecosystems, concentrated control, and higher misuse risk

The divergence between these two blocs will introduce significant operational and security challenges for companies operating across bloc lines. Global firms will find it difficult to maintain common assurance frameworks as architecture and deployment practices bifurcate. Regulators will likely struggle to benchmark systems built and governed under incompatible assumptions regarding transparency and responsibility.

“The U.S. is seeking AGI and ASI... They are trying to solve the whole problem [at once]... while China was adopting a more practical approach, using AI to solve real-world problems in development.”




- Ren Zhengfei
Founder, Huawei

Bottom line:

By the end of 2026, the strategic question will no longer be whether open or closed models prevail, but how this dual ecosystem shapes global access and norm-setting. The U.S. high-assurance model is likely to dominate premium markets where control and compliance are paramount, while the Chinese company-led open-weight ecosystem will foster rapid localization and experimentation across the developing world.

“As made clear in the President's Action Plan, the United States is committed to supporting the development and deployment of open-source and open-weight AI models...”

- Michael Kratsios
Director, Office of Science
and Technology Policy




China's Challenge Gets Complicated: Alternative AI Stack Not Fully Ready

The Chinese technological ecosystem is expected to reach a critical inflection point as innovation by leading Chinese companies furthers progress towards the creation of a largely domestic AI stack. This still-emerging sovereign infrastructure, centered on Huawei's Ascend processors and accompanying software ecosystem, will be further bolstered by advanced hardware and software offerings from emerging players such as Moore Threads, Meta, and Biren. Blockbuster IPOs from these firms and smaller but capable AI developers such as Zhipu (Knowledge Atlas), Minimax, and Moonshot in late 2025 and more in early 2026 provide major infusions of capital necessary to scale domestic capabilities. While this nascent ecosystem may not yet match the sheer efficiency of the Nvidia-led hardware environment, by late 2026 portions of it are likely to become "good enough" for training frontier-class models and support complex agentic applications.

In 2026, Chinese labs will run a dual-track procurement strategy. ByteDance, Alibaba, and DeepSeek are expected to receive significant volumes of Nvidia H200 class GPUs by early- to mid-2026 for the most demanding training, while expanding use of Huawei based systems for inference. This hedges against geopolitics while accelerating the buildout of a domestic software ecosystem that should reach greater depth by late 2027.

The Chinese AI stack progresses



Source: DGA-ASG Technology Policy Practice

China can train frontier-class models on domestic hardware, but pays a sovereign premium in efficiency and cost. High Bandwidth Memory (HBM) and advanced-node fab capacity remain the most binding constraints through 2026.

Trend: A Huawei ecosystem gains traction but Nvidia still preferred for training

Huawei is shifting from selling discrete GPUs to delivering full system platforms. In 2026, Ascend accelerators integrated into large scale clusters will allow some Chinese labs to complete major training runs entirely on domestic hardware, even as Nvidia systems remain in use for select workloads. Smaller but technically strong Chinese AI hardware firms will also begin winning niche deployments.

As domestic semiconductor capacity ramps late in 2026, Huawei will be better positioned to export a China-centric AI stack to non-aligned and Global South markets. This is already visible in Africa, where Huawei is pairing DeepSeek's open weight models with domestic hardware to deliver low-cost AI-in-a-box solutions to telecom and cloud providers. That model is likely to spread among countries seeking to avoid U.S. technology or sanctions, driving de facto standardization on a Chinese stack across parts of the Global South.

Challenge: Fabrication, memory, and tooling remain structural choke points

Despite these advances, the Chinese companies building the domestic AI stack will continue to face persistent structural challenges in 2026. Advanced semiconductor production remains heavily dependent on domestic foundry leader SMIC, whose 7-nanometer (nm) and prospective 5-nm manufacturing capacity is limited by U.S. export controls. Domestic memory leaders CXMT and YMTC are expected to make significant progress in producing more advanced memory, including high-bandwidth memory, a critical component of advanced AI hardware.

Blockbuster IPOs for both firms in 2026 will boost their capex spending and boost capacity. Later in 2026, additional capacity for advanced AI semiconductor production is expected to come online, benefiting a cohort of capable AI hardware startups that have or are expected to complete IPOs in the first half of the year.

“U.S. computing power is probably one to two orders of magnitude larger than China's. However, I've noticed that both OpenAI and Anthropic are investing a large portion of their computing power into next-generation research. We, on the other hand, are relatively strapped for cash; delivery alone likely consumes the majority of our computing resources. This is a significant difference.”

- Justin Lin
Technical Lead, Tongyi Qwen

Challenge: Multinationals must maintain two incompatible pipelines as China locks in the Ascend ecosystem

The lack of standards around agentic AI deployment at the enterprise level complicates cross-border data and operational issues.

Multinational corporations will find themselves forced to maintain two potentially incompatible technical pipelines. This institutional bifurcation will manifest in China-specific architectures and the ring-fencing of local operations. Over time, this technical divergence will make it more difficult for firms to deploy unified global AI solutions, unless a global standards process around agentic AI — launched in late 2025 under the Agentic AI Foundation — gains traction and includes leading Chinese AI developers.

Bottom line:

By leveraging a mix of legacy Western hardware for frontier research, while continuing to invest capital and resources into more capable domestic alternatives, Chinese firms will continue to straddle two digital worlds. They will contribute more effort than in 2025 to the development of a more China-centric AI stack. Firms that successfully navigate this split by building adaptable, dual-track operational pipelines will maintain their presence across critical markets.

8


Brussels Effect Hits AI Wall: Europe Aspires to be Major Player

The European Union is expected to attempt a strategic pivot away from heavy-handed regulation as it seeks to build a more competitive domestic industry. While the "Brussels Effect" — the EU's ability to set de facto global standards for privacy and digital markets — has been a defining feature of its global influence, AI is likely to mark the point at which that influence reaches a natural limit. Heavy early-stage rules, combined with Europe's weaker position in frontier compute, will make it increasingly difficult for Brussels to dictate how AI is built and deployed elsewhere.

Regional hubs in the Gulf, India, and Asia are likely to adopt only the regulatory components that suit their local interests while ignoring the rest. Simultaneously, European governments will grapple with the tension between strict precautions and fears of being left behind in the global industrial race. As implementation of the EU AI Act moves from principle to practice, the Commission must balance its role as a global rule maker with the risk that heavy requirements will push development and deployment to other regions.

At the same time, Brussels and major member states are moving to pair regulation with industrial muscle. The Commission is opening EuroHPC supercomputers to startups for model training and inference and building an EU-level sovereign cloud to keep sensitive workloads onshore, while France is using France 2030 and defense procurement to anchor firms like Mistral and Hugging Face, and Germany is scaling public compute and industrial AI pilots through its national AI and data center programs.

EU AI Act: Implementation Timeline (2025-2026)






Source: European Union, European Parliament

Challenge: Rules move faster than infrastructure

Despite these headwinds, Europe has an opportunity to define a distinct model of AI deployment centered on safety and verifiability. If the EU can successfully align AI Act implementation with real-world infrastructure and clear technical standards, it could give European firms a competitive edge in highly regulated sectors.

Bottom line:

The central question for 2026 is whether Europe can close the gap between its ambitions as a global regulator and its position as a late mover in compute and deployment. If implementation of the AI Act remains disconnected from industrial reality, the Brussels Effect may stall just as other global regions begin to assert their own rules.



9

U.S. Regulatory Efforts Hit Quagmire: Federal and State Tensions Create More Confusion

The U.S. is moving deeper into regulatory fragmentation throughout 2026 as states accelerate implementation of divergent AI frameworks while Washington remains unable to pass national legislation. The real U.S. regulatory baseline is now being set by states rather than federal lawmakers. The Trump administration has shifted from light-touch signaling to an explicit preemption strategy — establishing an AI Litigation Task Force, directing the Department of Commerce to identify "onerous" state laws, and threatening to condition federal funds on state policy choices — but these steps function more as uncertainty multipliers than as a clean reset of state authority.

Three overlapping state tracks are now defining U.S. AI governance: frontier model developer transparency and incident reporting such as CA SB 53 and NY RAISE Act; high-risk system governance focused on discrimination and consequential decisions, such as CO SB 24-205; and synthetic media rules that are proliferating rapidly with inconsistent definitions. California and New York have converged on a frontier developer framework requiring standardized safety protocols, internal governance structures, and incident reporting. Colorado's law embeds a "reasonable care" standard with defined documentation expectations and attorney general enforcement. These frameworks will drive early test cases and create the operational reality for firms in 2026.

Sample of AI-focused U.S. state bills

State	Bill	Status	Policy area	Who is regulated
 California	SB 53	Signed 2025; effective 2026	Frontier model transparency and incident reporting	Large frontier AI model developers
 Colorado	SB 24-205	Enacted 2024 still operative	Algorithmic discrimination	High risk AI system deployers
 Florida	SB 1344	Pushed in 2025	Child safety consumer AI	AI chatbot developers
 Missouri	HB 1747	Pushed in 2025	AI content labeling	Individuals, platforms, and creators.
	HB 2239	In pipeline	Data center transparency	Data center operators over 100 MW
 New York	S6953B / A6453B (RAISE Act)	Signed 2025; effective 2026	Frontier model safety, governance, and incident reporting	Large frontier AI model developers
	AB 8962	Pushed in 2025	Media transparency	News outlets
	AB 8595	In pipeline	Model training transparency	AI model developers
	AB 1456	In pipeline	Healthcare coverage decisions	Health insurers and model providers
	SB 8623	In pipeline	Algorithmic pricing	Businesses using pricing software
 Rhode Island	SB 2010	Pushed in 2025	Healthcare AI transparency	Health insurers and AI developers
 Tennessee	SB 1493	In pipeline	Child safety companion AI	AI developers
 Washington	SB 5956	Pushed in 2025	Child safety education	Schools

Source: DGA-ASG Technology Policy Practice

“...we are just two weeks into 2026. And yet the volume and complexity of state AI laws is at an all-time high. Many, if not nearly all, of these laws have extraterritorial effect. Almost all of them have gaps in drafting so large as to make any sane reader question whether the drafters really understand what they are doing.”

- Dean Ball
Lead Drafter of the U.S. AI Action Plan

Trend: U.S. firms face rising friction as state rules collide with EU obligations

This state-led experimentation has introduced deep uncertainty for American companies already grappling with the complexities of the EU AI Act. Throughout 2025, the consequences of European regulation became clear as major developers delayed product rollouts to navigate compliance timelines. The emergence of varying state-level requirements for watermarking, safety evaluations, and incident reporting will further complicate these deployment decisions. Conflicts between state rules and international standards will likely push multinationals to maintain overlapping compliance systems, significantly raising the cost of nationwide rollouts and risks stifling domestic innovation.

Challenge: Fragmentation raises operational and legal exposure

Fragmentation is an enforcement and litigation risk story. Colorado's "reasonable care" standard and California-New York incident reporting expectations increase the likelihood that near misses become regulator-facing issues driving investigations and follow-on claims. Synthetic media laws spread with inconsistent carve-outs, raising accidental noncompliance risks for national platforms. The result is higher baseline legal exposure where plaintiffs and state attorneys general can test novel liability theories in the vacuum of federal clarity.

Challenge: Federal attempts at harmonization will falter

Despite the White House's aggressive posture, viable federal harmonization remains unlikely. The Dec. 11 executive order titled Ensuring a National Policy Framework for AI, is explicit, but executive action does not automatically preempt state statutes and will face judicial constraints. In July 2025, the Senate voted 99-1 to strip a proposed 10-year moratorium on state AI regulation from a Trump-backed budget package. With Congress deadlocked, voluntary federal frameworks like National Institute of Standards and Technology's AI Risk Management Framework will remain influential for governance posture but nonbinding as legal shield.

Opportunity: Early movers can shape the eventual federal baseline

Fragmentation creates first-mover advantage for firms treating state compliance as an operating model.

California-New York frontier-developer transparency can become a de facto national benchmark before Congress acts. Colorado-style "reasonable care plus documentation" may template high-risk deployment governance for other states. Companies can operationalize a single internal baseline satisfying the strictest state requirements, reducing marginal rollout costs. Compliance ecosystems will expand around jurisdictional mapping, incident readiness, and audit-grade documentation.

“Investment in AI is helping to make the U.S. Economy the 'HOTTEST' in the World — But overregulation by the States is threatening to undermine this Growth Engine...”

- Donald J. Trump


Bottom line:

The U.S. remains stuck in a regulatory patchwork defined by state activism and federal paralysis. Companies that invest in adaptable governance systems and engage directly with state regulators will be best positioned to operate across this landscape. Conversely, firms that wait for federal clarity will face rising compliance costs and significant legal uncertainty as the domestic environment continues to bifurcate.

Advanced AI Diffusion Raises Risks: AI-Driven Incident More Likely

The rapid proliferation of frontier models suggests that 2026 will be the year AI-mediated harm transitions from a theoretical concern to tangible operational risk. As models gain autonomous reasoning, complex planning capabilities, and access to the physical world through agentic and embodied AI, the likelihood of a high-profile incident will refocus government attention on safety and security. While 2025 was defined by a reluctance to regulate for fear of stifling innovation, a major event in 2026 could serve as a catalyst for genuine international cooperation, potentially bridging the current divide between Washington and Beijing. As intelligence agencies increasingly attribute organized crime or state-sponsored activities to sophisticated AI tools, the likelihood of a serious incident causing loss of life will move from a tail risk to a recurring feature of the security landscape.

AI-Mediated Risk Strategic Assessment: 2026 Outlook



Note: Arrows indicate the trajectory of risk through 2026
Source: DGA-ASG Technology Policy Practice

Trend: Political appetite grows for hard guardrails on advanced systems

Open letters and campaigns calling for a slowdown in AI development are expected to intensify in 2026 as models become more capable and labs continue to signal progress toward artificial general intelligence (AGI). The "Statement on Superintelligence" released in late 2025 called for a prohibition on developing superintelligence until it can be shown to be controllable and safe, attracting a broad coalition of AI researchers and public figures. Widely discussed AGI/artificial superintelligence roadmaps, such as AI 2027, have brought the potential for catastrophic AI outcomes into mainstream discourse. In 2026, these efforts are likely to move toward concrete proposals for compute caps, binding evaluation regimes, and treaty-like arrangements on the most advanced systems. Early backing from a mix of researchers, lawmakers, regulators, and leading labs could turn these ideas into a live negotiating agenda.

Challenge: Hybrid human-AI incidents become the most plausible near-term threat

AI danger in 2026 is unlikely to come from a rogue autonomous agent. The more plausible pathway is a hybrid disaster in which human intent is amplified by AI. As models become more capable at software engineering and cyber-relevant tasks, smaller groups can attempt more cyber operations with fewer operators. Crises could be further exacerbated by agentic bots and synthetic media, which complicate attribution and undermine public trust in media and communications. The practical

risk extends to several cascading effects: delayed emergency response, misallocated resources, panic dynamics, and reactive policy moves that reshape the operating environment for firms deploying advanced systems.

Challenge: Biosecurity, chemistry, and industrial control systems become higher-stakes test beds

The biological and chemical domains remain the highest-stakes areas for potential AI-enabled risks. The realistic near-term concern is less that models create weapons in isolation and more that they lower barriers at specific stages of malicious workflows such as idea generation, literature synthesis, and operations, while strict materials access and limited lab knowledge still constrain feasibility. Despite internal red-teaming, sophisticated jailbreaks may allow nonexperts to refine approaches. At the same time, the growing reliance on partially agentic systems for monitoring in logistics, energy, and other infrastructure increases the risk of accidents stemming from automation or misspecified objectives. In a more widespread diffusion environment, these safety engineering risks can become as dangerous as deliberate misuse.

Opportunity: Early movers can shape incident response norms and future regulation

The emerging risk environment favors early movers across the organizational spectrum. Governments that build independent evaluation capacity will shape oversight regimes before any rules/laws are set. Insurers are moving beyond pricing to

active risk management, bundling cyber coverage with governance assessments and pre-breach services. Developers that transparently publish capability thresholds and document escalating safeguards under responsible scaling frameworks will gain regulatory credibility and differentiate themselves, especially if AI-incidents occur.




Bottom line:

In 2026, AI-mediated harm is likely to become more visible as agentic systems are deployed across sensitive domains. The first major incident, depending on severity, could trigger a rapid shift in regulatory

expectations, pushing the industry toward a more predictable compliance and liability environment. Such an event would also jumpstart stalled global discussions on creating a safety framework that is widely accepted and implemented in a verifiable manner.

“The best hope for a good future is a 'Chernobyl-scale AI disaster.'”

- Stuart Russell, recounting an AI CEO




About Us

DGA Group is a global advisory firm that helps clients protect – and grow – what they have built in today's complex business environment. We understand the challenges and opportunities of an increasingly regulated and interconnected world. Leveraging the experience and expertise of **Albright Stonebridge Group**, a leader in global strategy and commercial diplomacy, and a deep bench of communications, public affairs, and government relations consultants, we help clients navigate and shape global policy, reputational and financial issues. To learn more, visit dgagroup.com.



Please contact **Paul Triolo**, Partner and Technology Policy Lead, at Paul.Triolo@dgagroup.com, **Alexis Serfaty**, Associate Partner, at Alexis.Serfaty@dgagroup.com, and **Nikta Khani**, Senior Associate, at Nikta.Khani@dgagroup.com, with any questions or to arrange a follow up conversation.



Global coverage for global needs.
13 offices | 46 extended locations