

Two Models, One Race

How the United States and China are building AI, and where the deals are

A Lombard Global white paper · Part 1 of 4

By Bill Billeaud and Shrikant Patil · Lombard Global · July 2026

In brief

Most conversations about artificial intelligence mention only two countries: the United States and China. That is fair, they dominate, but it hides the more useful question: how are they winning, and which approach is the rest of the world copying?

The two superpowers are running the race in opposite ways. The United States is building a closed, proprietary stack: the best models stay behind an API, access is sold, and export controls keep the most advanced chips at home. China is making the opposite bet, open-weight and free to download, giving its models away to win global adoption it cannot yet buy with raw computing power.

That single split, closed versus open, shapes everything downstream: where the money goes, who buys chips and software from whom, and which companies become cross-border acquisition targets. This paper maps both countries side by side and draws out what the divide means for technology M&A. It is Part 1 of a four-part Lombard Global series on the global AI race, published weekly: Part 2 covers the Middle East and Europe, Part 3 covers India and the emerging markets, and Part 4 steps back to ask where the global AI race is headed and what it means for dealmakers.

The United States and China at a glance

Dimension	United States	China
AI “model”	Proprietary / closed: sell access to the best models	Open-weight / open-source: give models away
Leading labs	OpenAI, Anthropic, Google (closed); Meta Llama, OpenAI OSS (open)	DeepSeek, Alibaba Qwen, Moonshot Kimi, Zhipu GLM
Government posture	Light-touch and pro-build; move to preempt state law; managed chip exports	State-directed industrial policy; push for self-sufficiency
Data centers	Most of any country (~5,400 by Cloudscene’s count)	Large state-directed build-out (“Eastern Data, Western Computing”)
Flagship compute project	Stargate: \$500B / ~7 GW planned	National AI grid: ~\$295B drafted, 80%+ domestic chips
2026 build spend	~\$700-725B big-tech capex (~75% AI)	~\$70B provider data-center spend; cheap-energy edge
Private AI investment (2024)	\$109B, roughly 12x China	~\$9B
Venture funding (Q1 2026)	~\$250B (~83% of global)	~\$16B (state-led)
AI chips	Nvidia leads the world	Huawei Ascend / SMIC; memory and software gaps
Most valuable AI company	Anthropic ~\$965B; OpenAI ~\$852B	Private, state-linked, far smaller marks
Cross-border M&A	Open and acquisitive	Effectively closed in both directions

The United States: the closed frontier

The short version. The U.S. is not winning on one front; it is winning on four at once: policy, compute, capital, and deals. It hosts the most data centers, the most private AI investment, the most valuable AI companies on earth, and the most active AI M&A market. Its approach is proprietary-first: keep the frontier closed, sell access, and use export controls to keep the best chips at home. The one crack in that wall is open-weight models, where China sets the pace and the American answer is mostly Meta’s Llama and OpenAI’s open-weight releases.

Policy: build fast, stay closed. Washington has spent eighteen months turning “win the AI race” into formal policy: Executive Order 14179 (January 2025), the *Winning the Race: America’s AI Action Plan* roadmap of roughly 90 federal actions (July 2025), and a March 2026 National Policy Framework that asks Congress to preempt a patchwork of state AI laws in favor of one lighter federal standard ([Sidley](#); [White House](#); [Holland & Knight](#)). On chips, the U.S. shifted from hard bans to managed access: a January 2026 Commerce rule lets Nvidia sell H200s to China under a 25% tariff and case-by-case licensing. The logic is simple: sell the silicon, keep the very best at home, and tax the rest ([Council on Foreign Relations](#)).

Compute: pouring concrete. The U.S. hosts more data centers than any other country, roughly 5,400 by Cloudscene's widely-cited count, though tallies run lower on stricter definitions ([Visual Capitalist / Statista](#)). The Stargate venture (OpenAI, SoftBank, Oracle, MGX) is committing \$500 billion over four years, with close to 7 GW of capacity planned and the flagship Abilene, Texas campus already live ([OpenAI](#)). Across big tech, Amazon, Microsoft, Google, and Meta are guided to spend on the order of \$700-725 billion of capex in 2026, up sharply from roughly \$410 billion in 2025, with about three-quarters going to AI ([Tom's Hardware](#); [CNBC](#)). No other country is deploying private capital into compute at this scale.

Capital: where the money is. In 2024, U.S. private AI investment hit \$109.1 billion, about twelve times China's \$9.3 billion ([Stanford HAI](#)); in Q1 2026, U.S. companies raised roughly \$250 billion, about 83% of all global venture funding ([Crunchbase](#)). That capital produced the two most valuable private companies in the world: OpenAI at an \$852 billion valuation (a record \$122B round, March 2026) and Anthropic at roughly \$965 billion (May 2026), both moving toward IPOs ([CNBC, OpenAI](#); [CNBC, Anthropic](#)).

Deals: the agentic squeeze. The largest U.S. enterprise-software players are buying two things in 2025-26: AI agents and the data plumbing that feeds them. Salesforce paid roughly \$8 billion for Informatica (data management) and agreed to acquire the AI agent company Fin for \$3.6 billion; ServiceNow bought the agentic-AI firm Moveworks for \$2.85 billion and moved into data governance ([CIO Dive](#); [CMSWire](#); [ServiceNow 8-K](#)). The pattern is consistent: platforms and labs are paying premium prices to own agent capability and clean, governed data.

China: the open challenger

The short version. China cannot win the way America wins, so it is winning differently. Cut off from the best Nvidia chips, its labs made their models open, free to download and adapt, and gave them to the world. That open-weight strategy, plus the cheapest electricity of any major economy and a state willing to fund compute directly, has made China the clear number two and the supplier of choice for any country that cannot afford the American stack.

The open-weight bet. Where the U.S. keeps its best models closed, China gives its away. A year after DeepSeek's R1 model jolted the market, the country's top labs (DeepSeek, Alibaba Qwen, Moonshot Kimi, and Zhipu GLM) are shipping open-weight models at a furious pace, holding roughly four of the top five open-weight positions, with Qwen among the most-downloaded open-model families in the world ([CNBC](#)). A U.S. congressional review calls it a "two-loop" strategy: give models away to win global developers, then feed that scale back into China's industrial base ([USCC](#)). It is also a workaround: blocked from the best chips, Chinese labs competed on software efficiency and used openness to buy mindshare.

Compute: cheap power, state checkbook. China's edge is energy and coordination. It has the cheapest, most abundant electricity of any major economy, increasingly the binding

constraint on AI, and Beijing is spending directly to convert it into compute ([AI Jazeera](#)). The headline move is a drafted ~\$295 billion plan for a nationwide AI computing grid running on roughly 80% domestic chips, which structurally locks out Nvidia and AMD, built by state telecoms and targeted at around 2028 ([Bloomberg](#); [Tom's Hardware](#)). The catch: much of it is uncoordinated, with local-government competition producing overbuilt, idle capacity ([ASPI](#)).

Chips: closing the gap, stuck on memory. The whole strategy rests on domestic silicon. Huawei's Ascend line, made by SMIC, is the centerpiece, with an estimated 812,000 chips shipped in 2025 and projected AI-processor revenue near \$12 billion in 2026 ([Tom's Hardware](#)). Two constraints remain: high-bandwidth memory is the bottleneck, and Huawei has no real answer to Nvidia's CUDA software ecosystem. China is catching up faster than skeptics expected, but it is not self-sufficient yet.

Capital: state-led, not venture-led. On private capital, China does not come close to the U.S.: roughly \$16 billion of venture funding in Q1 2026 against America's \$250 billion ([Crunchbase](#)). The difference is who writes the checks: in China it is government guidance funds, state banks, and state-owned enterprises directing money to strategic priorities like chips and infrastructure ([Stanford FSI](#)). China competes on state direction and cost, not on the depth of its private markets.

A closed M&A loop. For dealmakers, China is notable mostly for the deals that cannot happen. U.S. capital is now restricted from funding Chinese AI by Treasury's Outbound Investment rule (effective January 2025) and the COINS Act (signed December 2025); CFIUS blocks Chinese acquisitions of U.S. technology; and China runs its own national-security review of outbound deals ([Holland & Knight](#); [Skadden](#)). China's AI consolidation therefore happens inside a domestic loop, largely closed to Western buyers. The cross-border action China creates is not at home; it is in the third countries it courts with free models and financed infrastructure.

What the divide means for cross-border dealmaking

Step back and the picture is simple. There are two models: the American (proprietary, closed, sold by the API and protected by export controls) and the Chinese (open-weight, given away to win the world). Almost every other country is quietly choosing a side or hedging between them. Three implications stand out for technology M&A:

1. "Whose model?" is now a diligence question. A target's choice between the closed U.S. stack and open Chinese models, or between Nvidia and Huawei silicon, shapes its chip access, customer base, and regulatory exposure. That alignment can raise or lower a company's deal-ability, and acquirers are starting to price it in.

2. The buyers are concentrated; the targets are dispersed. The most acquisitive players are U.S. enterprise-software strategics and, increasingly, Gulf capital. The most attractive targets, agentic-AI and data-engineering firms, are scattered across the U.S., Europe, India, and the emerging hubs. That gap is the cross-border opportunity.

3. The agentic squeeze is global. As AI agents compress headcount-based services revenue, mid-market IT-services and data-engineering firms everywhere face the same choice: retool, partner, or sell. A Western front end paired with lower-cost offshore delivery is fast becoming one of the most deal-able profiles in the market.

The US-China divide is not just a geopolitical story. It is a map of where capital is pooling, where targets are forming, and which alignments make a company easier or harder to sell. That map is what the rest of this series will fill in.

Sources

United States

- [Sidley, America's AI Action Plan and EO 14179](#)
- [White House, America's AI Action Plan \(PDF\)](#)
- [Holland & Knight, National AI Policy Framework \(March 2026\)](#)
- [Council on Foreign Relations, H200 export policy](#)
- [Visual Capitalist / Statista, data centers by country](#)
- [OpenAI, Stargate sites](#)
- [Tom's Hardware, big-tech 2026 capex](#)
- [CNBC, big-tech AI spend \(February 2026\)](#)
- [Stanford HAI, 2025 AI Index \(economy\)](#)
- [Crunchbase, Q1 2026 global venture funding](#)
- [CNBC, OpenAI \\$122B round at \\$852B](#)
- [CNBC, Anthropic tops OpenAI at \\$965B](#)
- [CIO Dive, Salesforce and Informatica](#)
- [CMSWire, Salesforce and Fin](#)
- [ServiceNow, FY2026 8-K](#)

China

- [CNBC, Chinese labs accelerate after DeepSeek](#)
- [USCC, "Two Loops" \(China's open-AI strategy\)](#)
- [Al Jazeera, China's cheap-energy edge](#)

- Bloomberg, China's ~\$295B AI build-out
- Tom's Hardware, China's \$295B AI grid
- ASPI, China's overbuilt compute
- Tom's Hardware, Huawei and SMIC Ascend ramp
- Stanford FSI, government venture capital and AI in China
- Holland & Knight, Outbound Investment Screening Rule
- Skadden, Treasury "reverse CFIUS" and the COINS Act

A note on sources: figures here are drawn from public company releases, government documents, and press reporting current to mid-2026. Where valuations or deals are private or fast-moving, we cite the most authoritative public source available.