




CATEGORY ROLLUP

Policy, geopolitics & China - 2026-W40

Week of September 28, 2026

The central narrative of this week's developments is a profound, structural shift in the AI competition: the transition from a race defined by software and algorithmic brilliance to one defined by the "materiality" of in

AI News Evolved · cognitoriverdelta.ai

CATEGORY ROLLUP

Policy, geopolitics & China - 2026-W40

Week of September 28, 2026 · 11 min read

Policy, geopolitics & China

Grid, gas & utilities

Policy, geopolitics & China · week 2026-W40: Sep 27 - Oct 03, 2026 · 1 subtopic(s) covered · 2526 words · expanded

Overview

The central narrative of this week's developments is a profound, structural shift in the AI competition: the transition from a race defined by software and algorithmic brilliance to one defined by the "materiality" of intelligence—the hard, physical constraints of electricity, silicon, and specialized infrastructure. While the previous year was dominated by debates over Large Language Model (LLM) capabilities and "alignment," the focus has now pivoted toward the physical scaffolding required to sustain them. We are witnessing the emergence of a "Compute-Energy-Geopolitics" triangle, where national security is no longer just about who writes the best code, but who can power the largest data centers and secure the most sophisticated lithography.

This pivot is driving a wedge between two fundamentally different models of development. On one side is the United States, characterized by a high-innovation but uncoordinated and laissez-faire approach, currently struggling with domestic friction regarding power grid capacity and public skepticism. On the other is China's state-led, highly structured model, which treats AI as a core component of national industrial policy, integrated through five-year plans and massive, state-directed energy and computing projects.

The week's threads suggest that the "AI race" is no longer a sprint toward a singular intelligence threshold, but a marathon of resource mobilization. As US policy shifts from regulating model behavior to controlling the "physicality" of AI—specifically GPU allocation, data center access, and energy supply—the geopolitical friction points are moving from the

digital realm (the internet and software) to the physical realm (export controls, investment bans, and the electrical grid).

AI Policy

The landscape of AI regulation is undergoing a fundamental metamorphosis. For much of the recent past, policy discussions have centered on "model governance"—the oversight of individual frontier models to prevent misuse, such as the creation of biological weapons or cyber-attacks. However, as predicted by Jason Calacanis, there is a decisive regulatory pivot underway. Within the next 12 to 18 months, the focus of oversight is expected to shift from the software itself toward the control of the physical infrastructure required to run it: compute, data center access, and GPU allocation. This shift is being accelerated by the signing of the White House Accord on Super Intelligence, which signals a move toward treating high-end compute as a matter of critical national security rather than just a commercial enterprise. This transition fundamentally changes the regulatory burden; instead of auditing code, the state will increasingly audit hardware consumption and infrastructure access.

This transition brings a new layer of complexity to the concept of "safety." While regulators are looking at hardware, safety advocates are looking at the recursive nature of the technology itself. Dario Amodei has proposed a multi-layered safety framework that seeks to move beyond simple model testing. His vision includes the external evaluation of frontier labs and the establishment of "coordinated speed limits" among democratic nations to prevent a race to the bottom. Crucially, Amodei argues for the necessity of safety coordination with authoritarian governments to mitigate catastrophic risks. This highlights a growing tension in policy: the desire to maintain a competitive edge against rivals like China via technological denial versus the necessity of global safety guardrails to prevent unintended, recursive escalation.

In China, AI policy remains deeply intertwined with ideological control and state stability, creating a regime where technical capability and political alignment are inseparable. The Cyberspace Administration of China (CAC) continues to enforce a regime that ensures generative AI adheres to "Socialist Core Values." This is not merely a content moderation effort; it is a deep structural requirement that includes mandatory security assessments and the registration of all algorithms through the CAC's public filing system, the *Internet Information Service Algorithmic Recommendation Filing*. Furthermore, China is setting rigorous technical standards for "deep synthesis," requiring mandatory watermarking and metadata tagging on all AI-generated audio, video, and text content. This creates a bifurcated

policy world: the West is increasingly focused on the *safety of the output and the security of the supply chain*, while China is focused on the *alignment of the output with state ideology and the integration of the technology into the national industrial fabric*.

Geopolitics

The geopolitical dimension of AI has become a high-stakes game of "containment versus integration." The United States is doubling down on a strategy of technological denial, primarily through the Bureau of Industry and Security (BIS). The evolution of these export controls—moving from simple bandwidth caps to complex metrics like Total Processing Performance (TPP) and Performance Density—demonstrates a sophisticated attempt to close the loopholes that allow for "chiplet-based" workarounds. The US is not just targeting individual chips; it is targeting the entire ecosystem, including the lithography systems required to make them (such as ASML's EUV and DUV variants) and the very capital that fuels Chinese innovation, through the Treasury Department's outbound investment controls (Executive Order 14105). This strategy aims to starve the PRC of the high-end compute necessary for frontier training and military modernization.

This containment strategy is meeting a highly organized counter-strategy in China. Rather than attempting to build a single "super-model," China is building a parallel, state-directed computing ecosystem. The "AI+" initiative, announced during the 2024 National People's Congress, aims to weave LLMs into the very marrow of Chinese industry, from manufacturing to energy to governance. To bypass the "compute gap" created by US restrictions, China is deploying the "Eastern Data, Western Computing" (*Dongshu Xisuan*) project. This massive undertaking coordinates data centers in western provinces like Guizhou and Ningxia, pooling domestic hardware—such as the Huawei Ascend 910 series and Cambricon MLU accelerators—to create a unified national compute network that maximizes the utility of available domestic silicon.

On the multilateral stage, the tension is visible in the divergent paths of international governance. While the US successfully led a UN General Assembly resolution on AI in March 2024, China countered in July 2024 with a resolution focused on bridging the "digital divide" between the Global North and the Global South. This suggests that the geopolitical battle for AI influence is being fought not just in the laboratories of Silicon Valley or Beijing, but in the diplomatic corridors of the UN, where China is positioning itself as the champion of AI accessibility for developing nations, potentially offsetting the impact of US-led export controls by building a coalition of the "technologically underserved."

China

China's approach to AI is increasingly defined by its ability to synchronize technological ambition with physical capacity. Amy Webb identifies a stark divergence here: while the U.S. operates through a fragmented, laissez-faire model, China utilizes its five-year plans to create a highly coordinated, state-led roadmap for R&D and adoption. This is not just about funding; it is about the strategic, long-term allocation of the nation's physical resources, including broadband and infrastructure, to ensure that technological advancement is a state-directed outcome rather than a market-driven accident.

The most critical component of this strategy is the management of energy and infrastructure. David Carbutt highlights a massive disparity in the scaling of power capacity: China is currently doubling its power capacity every decade, whereas the United States has seen only 16% growth over the last 25 years. This disparity in "energy velocity" is a fundamental component of the AI race. If AI is a game of scale, and scale is a function of electricity, then China's state-backed expansion of its electrical grid provides a massive, long-term structural advantage. This capacity allows China to sustain the massive energy demands of its "Eastern Data, Western Computing" hubs and its broader "AI+" industrial integration.

This capacity is matched by a social environment that is, at present, more receptive to AI integration than its American counterpart. Carbutt notes a striking psychological divide: while roughly 80% of Americans express fear regarding AI, approximately 80% of Chinese citizens are pro-AI. This social consensus, combined with the state's ability to direct capital toward "AI+" industrial initiatives, suggests that China is building an environment where AI is treated as a public utility—a tool for national advancement—rather than a disruptive, feared force. This receptivity may allow for faster deployment of AI-driven industrial processes and a lower threshold for domestic infrastructure expansion.

Cross-cutting themes

The most significant cross-cutting theme this week is the **"Materialization of AI."** We are seeing the intersection of AI policy, geopolitics, and China through the single lens of physical resources.

- **Policy meets Geopolitics:** The regulation of AI is no longer just about the "code" (Policy); it is about the "compute" (Geopolitics). The US's move to regulate GPU allocation and data center access (as predicted by Calacanis) serves as the bridge

where domestic policy becomes a tool of international containment. By controlling the physical inputs, the US aims to enforce its geopolitical goal of technological denial.

- **Geopolitics meets China:** The "containment" strategy of the US (export controls and investment bans) is meeting a "structural" strategy from China (five-year plans and massive energy expansion). This creates a reactive loop: US hardware restrictions drive China to accelerate domestic hardware self-sufficiency (Huawei/Cambricon) and create massive, energy-intensive computing hubs in its western provinces to maximize the efficiency of their existing, non-US silicon.
- **The Energy-Security-Compute Nexus:** There is a growing tension between the *computational requirements* of AI and the *physical constraints* of the electrical grid. This is where the race could be won or lost. If the US cannot solve its grid bottlenecks and public opposition to infrastructure, its "uncoordinated" model may find itself unable to scale, regardless of how much software innovation it produces. This links domestic energy policy directly to the global AI security hierarchy.

Where sources agree

- **The Criticality of Infrastructure:** There is a consensus among analysts (Calacanis, Carbutt, Webb) that the AI competition is moving from a software-centric phase to a physical-infrastructure phase. The ability to secure compute, energy, and land is now viewed as a primary determinant of success, moving the goalposts from algorithmic development to resource mobilization.
- **Divergent Models of Governance:** Sources agree that the US and China are operating under fundamentally different systemic architectures. The US relies on a market-driven, uncoordinated, and laissez-faire model, whereas China utilizes a state-led, structured, and integrated model driven by five-year plans.
- **The Shift in US Policy Focus:** There is an emerging recognition that US regulatory and security efforts are pivoting from the oversight of model *outputs* (what the AI says) toward the control of the *inputs* (the hardware, the capital, and the energy).

Where sources disagree

- **The US "Laissez-faire" Model:** There is a tension regarding whether the US's uncoordinated approach is a strategic weakness or a strength. Amy Webb views the lack

of coordination as a major disadvantage compared to China's state-led roadmap. However, the recent moves toward the "White House Accord on Super Intelligence" suggest that the US government is attempting to "force" coordination through executive action to bridge this gap.

- **The Necessity of International Cooperation:** There is a significant philosophical divide regarding how to manage global risk. Dario Amodei advocates for a model that includes safety coordination with authoritarian governments to mitigate existential, recursive AI risks. This stands in direct tension with the US's broader geopolitical strategy of technological denial (BIS and Treasury controls) aimed at preventing China from reaching the same frontier capabilities.
- **The Future of Open-Weights Models:** There is an ongoing, unresolved debate in the policy community regarding the proliferation of open-weights models (such as Alibaba's Qwen or DeepSeek). The debate centers on whether these models act as a democratic advantage by democratizing access, or as a security vulnerability that necessitates stricter export controls to prevent dual-use capabilities from spreading.

Numbers and claims to verify

- **Power Capacity Growth:** Verify the claim that China doubles its power capacity every decade while the U.S. has grown only 16% in 25 years (Source: David Carbutt).
- **Public Sentiment Statistics:** Verify the claim that 80% of Americans fear AI while 80% of Chinese citizens are pro-AI (Source: David Carbutt).
- **Compute Thresholds:** Verify the specific TPP (Total Processing Performance) metric of $>4800\$$ and how it is applied to prevent chiplet-based workarounds in US export rules (Source: BIS framework).
- **The "White House Accord on Super Intelligence":** Verify the existence and specific regulatory implications of this accord as a driver for infrastructure-based regulation (Source: Jason Calacanis).

Investment and strategic implications

- **Infrastructure as the New Alpha:** For strategic planners and investors, the "AI play" is shifting away from pure-play software companies toward the "picks and shovels" of the physical layer. This includes energy providers, electrical grid technology, advanced

semiconductor manufacturing equipment (lithography), and specialized data center real estate.

- **The "Know Your Customer" (KYC) Mandate for Compute:** As the US Department of Commerce explores KYC requirements for Infrastructure-as-a-Service (IaaS) providers to prevent "cloud workarounds" in Southeast Asia and the Middle East, companies providing cloud compute will face significantly higher compliance, due diligence, and operational costs.
- **Geopolitical Risk in Capital Allocation:** The US Treasury's outbound investment controls (EO 14105) mean that private equity and venture capital must now treat any investment in Chinese "frontier" sectors (quantum, AI, semiconductors) as a high-compliance, high-risk activity that may be subject to federal prohibition.
- **Domestic Resource Mobilization:** In the US, the tension between AI development and the electrical grid suggests that companies with "on-site" power generation capabilities or those integrated into microgrid technologies may have a significant competitive advantage in the race to scale, as they can bypass traditional grid bottlenecks.

What to watch next week

- **Implementation of TPP Metrics:** Watch for any specific guidance from the Bureau of Industry and Security (BIS) on how "Total Processing Performance" will be audited for mid-tier chip manufacturers, as this will define the next generation of hardware trade legality.
- **Developments in "Eastern Data, Western Computing":** Look for announcements regarding the scale of new mega-datacenter hubs in China's western provinces (e.g., Guizhou or Ningxia), which would signal the progress of their domestic compute-pooling and hardware integration strategy.
- **Bilateral Safety Talks:** Monitor any follow-up communications from the Geneva-based US-China bilateral AI risk talks, particularly regarding "unintended escalation pathways" in nuclear command-and-control (NC3) and autonomous weapons limitations.
- **Energy Policy Debates:** Watch for domestic US legislative movement or regulatory filings related to the easing of grid interconnection bottlenecks for massive data center projects, as this is the primary physical constraint on US scaling.

Appendix: Individual perspectives

- **Amy Webb:** Characterizes the US approach as uncoordinated and laissez-faire, contrasting it with China's state-led strategy of utilizing five-year infrastructural and broadband plans to drive R&D and adoption.
- **Dario Amodei:** Advocates for a multi-layered AI safety approach including external evaluation of frontier labs, coordinated speed limits among democratic nations, and safety coordination with authoritarian governments to mitigate recursive AI risks.
- **David Carbutt:** Highlights a critical national security threat in the US stemming from public opposition to infrastructure and electrical grid bottlenecks, noting a massive disparity in power capacity growth between China and the US, as well as a significant divergence in public sentiment toward AI.
- **Jason Calacanis:** Predicts a major regulatory pivot occurring within the next 12 to 18 months, moving from the oversight of individual models to the control of physical infrastructure, data center access, and GPU allocation as matters of national security.

Related reports

CATEGORY ROLLUP

Grid, gas & utilities - 2026-W40

The defining storyline of this week is the intensifying "collision of timescales" between the exponential, near-instantaneous scaling of AI compute and the multi-decadal inertia of the global electrical grid. We are with

Week of September 28, 2026 · Grid, gas & utilities · Nuclear & SMRs

CATEGORY ROLLUP

Nuclear & SMRs - 2026-W40

The defining narrative of this week in the nuclear sector is the emergence of a high-stakes "AI-Nuclear Nexus." We are witnessing a fundamental and permanent shift in the identity of the nuclear consumer. For decades, th

Week of September 28, 2026 · Nuclear & SMRs · Grid, gas & utilities

DAILY AI BRIEFING

The White House Accord on Super Intelligence

The artificial intelligence industry is rapidly transitioning from model-centric development to infrastructure-driven, physical-world scaling, where regulatory focus is shifting toward data center access, energy grids, and national security. While industry leaders debate safety accords, compute econ

October 2, 2026 · Policy, geopolitics & China · Humanoid & general robotics

Everything here is free, and reader-supported.

Join free on Patreon to get the weekly digest by email, or become a member to get each morning's briefing by email before it reaches the site, request custom analysis, and get early access to new research tools. What members get

♥ [Join or support on Patreon](#)

Informational analysis synthesized by AI from sourced, dated material, curated by a human. Treat specific claims as unverified until checked. Not financial advice.

[About](#) · [How this is made](#) · [Corrections](#)