



SPECIAL COMPETITIVE
STUDIES PROJECT

OCTOBER 2026

The Open Frontier

How America Wins the AI Diffusion Race

A Letter from SCSP's President

The United States can lead at the technological frontier and still lose the most consequential part of the AI competition. That is the argument of this report, and it is why we wrote it.

Since the NSCAI, our work has rested on a simple point: innovation power is not only the ability to invent, but the ability to turn invention into adoption. We have seen the alternative. America led the world in wireless research, and Huawei still connected much of the globe by out-competing on cost, financing, and speed of deployment. Open models are where that contest is now being played. Most countries will never train a frontier model, but nearly all will deploy one. Their choice will help decide whose technology becomes the global default.

We are releasing this now because the evidence has moved. In January 2024, 1 percent of new fine-tunes published to Hugging Face were built on Alibaba's Qwen. By February 2026, the figure was 69 percent. Meta's Llama, the leading American open model, fell from a peak of 44 percent to 11. The most capable open models trail the leading closed systems by only about four months. Meanwhile, Washington too often discusses open models as though all of them are Chinese, and Chinese models as though all of them are unsafe. Neither follows from the other.

The window is still open, which is the other reason for the timing. In July, more than two hundred organizations signed an open letter making the case for American open-weight leadership. In August, the Department of Energy launched the Genesis Open Models Initiative. American developers are releasing large open models of their own. Defaults harden, though. Once developers have built on a model and wired it into their infrastructure, displacing it is expensive. The choices made in the next few years will matter more than those that follow.

On September 21, we convened a half-day workshop under the Chatham House Rule with companies, frontier model developers, government officials, and academics. This report draws on that discussion, though the analysis and recommendations are ours. We thank every participant for their candor.

The report sets out three lines of effort: build the ecosystem America can win, turn leadership into adoption, and govern by capability and use rather than by whether a model's weights are public. That last position is consistent with SCSP's view since Mid-Decade Challenges to National Competitiveness, and with what the NSCAI concluded five years ago about justified confidence in AI systems. The United States does not need to copy anyone's strategy. It needs a better one, and it needs to run it while the choice is still open.

This competition will not be decided by who builds the best model. It will be decided by whose models the world builds on. That can and should be the United States.



Ylli Bajraktari
President

Table of Contents

Executive Summary	3
China’s Open-Model Bet Is Working	4
The American Answer: Three Lines of Effort	7
Build the AI Ecosystem America Can Win	7
Turn Leadership into Adoption	11
Govern American AI Without Making It Inaccessible	14
The Metric That Matters: Who Is the World Building On?	18
Acknowledgements	20

Executive Summary¹

The United States can lead at the technological frontier and still lose the most consequential part of the AI competition, if developers, firms, and governments at home and abroad build on another country's models and infrastructure. Open models will help determine which technology ecosystem becomes the global default. Building on decades of open-source software development, openness in AI enables others to inspect, adapt, and improve shared technology, broadening participation and accelerating innovation. Openness spans the spectrum, from publicly available model weights to broader access to code, training data, and methods, with licenses determining how these resources can be used and shared. Open-weight models, systems whose learned parameters (i.e., the billions or even trillions of numerical settings a model adjusts during training to shape how it responds) are publicly released, allow users to download, evaluate, customize, and deploy them on infrastructure they control. History shows invention alone does not secure lasting leadership. For example, although the United States led the world in wireless research, Huawei won the diffusion race and connected much of the globe by out-competing on cost, financing, and speed of deployment.

China is implementing a similar strategy to AI. Its developers are using capable, low-cost, downloadable models to expand their footprint across the global AI ecosystem. The advantage does not come from openness itself, but the adoption it enables: lower costs for initial adoption as well as for switching models, the ability to run a model on one's own hardware, customization, a larger base of derivative models and tools, and a growing community with reasons to develop further within the same technical ecosystem. What matters about these models is not how they score on benchmarks, the standardized tests the industry uses to assess a model's capabilities. It is that China's open models are becoming the foundation on which developers are building precisely because they are open.

¹ The analysis and recommendations in this memo are based on an SCSP workshop that convened industry, academia, and government on September 21, 2026 under the Chatham House Rule.

In Washington, meanwhile, open models are discussed as though all of them are Chinese models, and Chinese models are discussed as though they all are unsafe. Neither follows from the other.

While there are trust and security concerns regarding Chinese open models, the American response should be to strengthen its own open model ecosystem alongside its proprietary frontier laboratories. The two ecosystems are structurally different. Leading U.S. frontier laboratories monetize access through subscriptions, enterprise products, and managed platforms; that commercial engine, rather than large public investment, funds the next generation of models. But a strong American strategy should also support open model developers and the researchers, startups, and businesses that build on their work. The United States needs a strategy that builds on its distinctive strengths: invest in a thriving open AI ecosystem by expanding access to compute, capital, and shared infrastructure; accelerate adoption of American AI across society; and govern risks based on what models can do, how they are used, and their impacts and outcomes, not simply whether their weights are public.

China's Open-Model Bet Is Working

China recognizes the strategic advantage of leading an open model tier. Most countries will never train frontier models, but nearly all will deploy them. Which should they adopt? Beijing seeks to frame their choice as inexpensive open Chinese models against expensive, closed American ones.

A capable model that is inexpensive to access and easy to run on one's own hardware lowers the barrier for developers to experiment, adapt it to their own needs, keep their data and intellectual property in-house, and build new systems on top of it.² Each successful deployment produces new applications, datasets, tooling, documentation, training expertise, and business relationships. Those complementary assets make the underlying model more useful, which makes the next developer more likely to choose it. That is the playbook China is running.


Developer adoption has shifted measurably toward Chinese open models. One useful indicator is Hugging Face, the main public repository for open AI models, where

² [The Strategic Implications of Open-Weight AI: Diffusion, Competition, and the Shifting Technological Frontier](#), Special Competitive Studies Project (2025).

developers publish fine-tunes and adaptations (models built by adjusting an existing one). In January 2024, 1 percent of those new models were built on Alibaba’s Qwen and by February 2026 jumped to 69 percent.³ Meta’s Llama, the leading American open model at the time, went the other way: it peaked at 44 percent in August 2024 and has since fallen to 11 percent.⁴ By February 2026, Chinese laboratories together accounted for roughly 70 percent of all new models built from an existing one.⁵ And on OpenRouter, a service that routes AI requests from applications to the model that answers them, Chinese open models grew from 2.8 percent of the requests in August 2024 to more than 70 percent by January 2026.⁶ And open models are running real products. One major commercial gateway, a service that businesses use to connect their applications to AI models, reported that open-weight models of all origins handled 56 percent of token volume in August 2026.⁷ These measures indicated that open models carry a growing share of use, and the most capable lag the leading closed systems by an average of only four months.⁸

Chinese open models have taken the developer base

Share of new fine-tunes and adaptations published to Hugging Face



Source: The ATOM Report: Measuring the Open Language Model Ecosystem (2026).
Each bar is an observed value. No intermediate points are implied.

³ [The ATOM Report: Measuring the Open Model Ecosystem](#), arXiv (2026).

⁴ [The ATOM Report: Measuring the Open Model Ecosystem](#), arXiv (2026).

⁵ [The ATOM Report: Measuring the Open Model Ecosystem](#) arXiv (2026).

⁶ [The ATOM Report: Measuring the Open Model Ecosystem](#) arXiv (2026). The OpenRouter figure measures the share of tokens served among open models only, drawn from the top ten open models per month.

⁷ Amelia Charles & Eric Dodds, [Open-weight models take 56% of token volume](#), AI Gateway Production Index, Vercel (2026).

⁸ Jack Edwards & Luke Emberson, [Open models lag state-of-the-art closed models by 4 months](#), Epoch AI (2026).

Chinese open models have displaced American ones in the developer base. Each bar is an observed value; no intermediate points are implied.

The releases behind those numbers show the pattern. DeepSeek publishes weights with instructions for running the model locally; Qwen has released weights across multiple sizes and formats;⁹ and Moonshot AI's Kimi family has released code and weights under permissive licenses, most recently Kimi K3 in July 2026, which at 2.8 trillion parameters (a measure of a model's size) is the largest open-weights model released to date.¹⁰

American industry increasingly recognizes what is at stake. In July 2026, a coalition of American technology companies asserted in an open letter, *Open Weights and American AI Leadership*, that U.S. leadership will be judged not by any single frontier model but by whether America builds an open ecosystem that diffuses into every sector of the world economy.¹¹ More than two hundred organizations had signed within a week.

The United States open model ecosystem has begun to take shape. In 2026, U.S. developers released or announced large open models, including Reflection AI's Beam, Thinking Machine's Inkling, and NVIDIA's Nemotron 3 Ultra.¹² Together, they show that American companies will be competitive with Chinese open-weight models. Shaping the American open model ecosystem requires government leadership as laid out in this memo. The diffusion race will reward the country that treats open and closed models as complements, not direct competitors.

⁹ Caiwei Chen, [What's next for Chinese open-source AI](#), MIT Technology Review (2026).

¹⁰ Daniel Butler, [Kimi K3 Open Weights Drop July 27: Near-Frontier Coding, Undisclosed Hallucination Risk](#), Tech Times (2026). Open release is also becoming a route to revenue. Alibaba plans revenue-sharing terms on the open release of its next Qwen flagship, triggered once a company earns more than \$20 million a year offering the model as a service; Moonshot AI already applies such a term to Kimi K3, at up to 30 percent. Reuters, reported August 7, 2026; see Cris Tolomia, [Alibaba is planning to charge big commercial users of its next open-source AI model](#), Quartz (2026).

¹¹ Nick Lichtenberg, Jensen Huang just used his first ever X post to warn the AI industry not to make the mistake that software narrowly avoided in the 1980s, Fortune (2026).

¹² [Introducing Beam: Reflection's 501B open-weight model](#), Reflection (last accessed 2026) (weights will release late October 2026); [Inkling, Open weights, ready to tinker](#), Thinking Machines (last accessed 2026); [NVIDIA Nemotron](#), NVIDIA Developer (last accessed 2026) (NVIDIA Nemtron is a family of open models with open weights, training data, and recipes).

The American Answer: Three Lines of Effort

Three lines of effort should organize the American open model strategy: build an AI ecosystem that plays to American strengths, adopt American AI across society, and govern risks according to capability, deployment context, and impacts and outcomes rather than by whether a model's weights are public.

Build the AI Ecosystem America Can Win

The rest of this strategy rests on one condition: that the United States remains the best place in the world to start, scale, and sustain AI companies across the stack. That depends on compute, capital, and shared infrastructure, and on a deliberate decision to support an open model ecosystem that researchers, startups, and governments can inspect, adapt, and trust.

1. Cost is a binding constraint on compute access

Compute, the specialized processors used to train and run AI models, is a cost barrier to an open model ecosystem far beyond frontier pretraining. American open model developers train models too, not quite at the largest frontier scale, but still requiring significant capacity. Startups and academic researchers need the capacity to run models, fine-tune them for particular uses, generate training data, and evaluate results. And government research programs and established industry partners also face rising costs and limited availability.

Demand from large commercial workloads has absorbed most available capacity, including older hardware that might otherwise support public research. Nor does a large pool of idle compute exist to be redistributed. If only the largest frontier laboratories can afford sustained access to compute, the American AI ecosystem will narrow precisely as foreign open-model ecosystems expand.

The federal response is the Genesis Mission, established by executive order in November 2025¹³ and led by the Department of Energy (DOE). Its purpose is broad: to accelerate American science by connecting federal supercomputers, experimental facilities, datasets, and AI systems, with priority domains including advanced

¹³ E.O. 14363, [Launching the Genesis Mission](#), The White House (2025).

manufacturing, biotechnology, critical materials, fusion and fission energy, quantum information science, and semiconductors.¹⁴

In August 2026, DOE launched the Genesis Open Models Initiative within that mission: a program to produce open-weight foundation models built for scientific discovery, in materials, energy systems, earth systems, fusion, biology, and high-energy physics. Open models are now a named part of Genesis. What the initiative lacks is appropriate scale to diffuse open models domestically.

Two distinct problems stand in the way. The first is funding the purchase of compute access.¹⁵ DOE holds the authority to purchase computing resources but lacks the appropriations to do so at scale; with more resources, it would direct more of them to open-model development. The second is authority to finance the construction of additional compute infrastructure through loans or loan guarantees. No agency clearly holds authority to finance buildout of much-needed compute infrastructure for open-model research.¹⁶ DOE's own financing authority covers energy, not compute.

Funding Genesis at scale also has a rationale beyond open models. Pooling federal demand produces more total compute than dividing the same resources among smaller clusters with separate power, networking, and facilities. Pooling works, however, only if access and contracting allow participating agencies to use what is built and if those agencies retain the budgets needed to consume it. Genesis is not a substitute for agency funding. It is a concrete program for federal research investment whose results Americans can see in medicine, energy, and materials, at a moment when public argument about AI is dominated by displacement and risk.

The idea is to fund the Genesis Mission to its full ambition. A narrower path runs through the National Nuclear Security Administration (NNSA), which also oversees three national laboratories. NNSA's programs are authorized each year in the National Defense Authorization Act, which passes reliably, and funded through the Energy and Water appropriations bill, where the laboratories have a durable

¹⁴ [Department of Energy Genesis Mission Home Page](#), U.S. Department of Energy (last accessed 2026).

¹⁵ Alfred Ng, [Tech firms pledge \\$2.4B inc computing power to Trump's Genesis Mission](#), Politico (2026).

¹⁶ The Department of War can finance infrastructure through its Office of Strategic Capital, but its investments are tied to specific military capabilities rather than to research.

constituency.¹⁷ NNSA activities include advancing computation: while second-best, funding through such an existing authority could support an open-model mission at DOE, in the absence of a fully-funded Genesis mission. However, such an approach risks attaching open-model research to a nuclear security mission rather than treating it as a national research priority, and it reaches only part of what a full program would cover. But it is still a meaningful step forward in leveraging open models because NNSA has significant needs for effective, efficient AI models for a variety of different uses. If that route is taken, a designated laboratory or laboratory consortium needs an explicit mandate and dedicated funding together. A laboratory handed the responsibility without the money will not deliver.

2. The process for accessing compute must be streamlined

The process for accessing compute must be easier. Reaching compute under DOE's Genesis Missions means clearing many separate process steps (e.g., a separate contract for each engagement and a separate security review for each system), each with its own requirements and its own timeline. A large contractor with an existing relationship has already cleared most of them. A three-person start-up faces all of them at once, and it is the cumulative burden, not any single rule, that keeps it out. The Department of War's DIU Bridge program already consolidates process steps like clearances and accreditation into a single pathway for small companies; Genesis has no equivalent.¹⁸ A company whose work directly relates to the mission should be able to make that case once, rather than clear a dozen separate hurdles to reach the same compute. The steps to reach Genesis Mission compute should be consolidated, and intake should be continuous. For qualified startups and small businesses, this should mean single application, coordinated contracting and security reviews, and affordable compute allocations. Fewer administrative steps alone will not address the cost of access.

¹⁷ [Energy and Water Development Appropriations: Nuclear Weapons Activities](#), Congressional Research Service (2018). "Congress authorizes funding for both DOD and NNSA nuclear weapons activities in the annual National Defense Authorization Act (NDAA). While Congress considers appropriations for DOD's nuclear weapons activities in the Defense Appropriations bill, it funds the NNSA budget through the Energy and Water Development Appropriations bill."

¹⁸ [Defense Innovation Unit, DIU Bridge Program](#), Defense Innovation Unit (last accessed 2026).

3. Support fully open models where it matters most

Open weights are not the same as a fully open model. Weights can be downloaded and fine-tuned. A fully open model also documents how it was built, including the training data, training methods, evaluation results, and materials needed for reproduction. Scientific researchers need that complete record because a model's output in domains like biology, materials science, or chemistry may make claims about physical systems that others must be able to verify. National security users need it to understand the provenance and limits of a model they deploy. Disclosure varies by model and should be assessed directly rather than inferred from country of origin. Open weight models can support scientific and security work, even where additional disclosures are necessary to establish data provenance or reproducibility.

Commercial open model developers face strong incentives not to provide that record. A company's training data and methods are among its most valuable IP, and publishing them hands critical IP to competitors. American companies have released permissively licensed model weights and supporting tools to encourage developers to build on their technology, but such releases do not necessarily make the training process reproducible.¹⁹ U.S. based research organizations have gone further by releasing training data and code.²⁰ These are significant contributions, but they do not by themselves ensure that fully open models will be developed and maintained in areas of greatest public need. If the United States wants fully open models in the fields where they matter most to Americans, it will have to fund their production through academic, nonprofit, and national laboratory programs, with full release as a condition of support.

Recommendations

1a. Fully resource the Genesis Mission so that it can benefit open model ecosystem development. DOE holds the purchasing authority and has launched the open models program; however, it lacks the dedicated resources to run either at scale. Alternatively, an open-model mission can be assigned to NNSA and funded through its

¹⁹ [Introducing MPT-7B: A new Standard for Open-Source, Commercially Usable LLMs](#), databricks (2023).

²⁰ [Language models, OLMo 2](#), Ai2 (last accessed 2026); [EleutherAI pythia](#), Github repository (last accessed 2026).

existing authorization and budgetary processes, provided any designated laboratory or laboratory consortium receives an explicit mandate and dedicated funding together.

1b. The Department of Energy should explore pathways to finance, or guarantee financing, for large-scale compute infrastructure dedicated to open-model research. DOE's existing credit and loan programs are limited to energy supply projects, and no federal agency has a dedicated financing authority for open model compute infrastructure.

2. Streamline access to Genesis Mission compute. Compute access should require a single justification, with consolidated process steps and rolling applications.

3. Fund fully open-source models for scientific fields critical to Americans and national security missions. Models that release training, data documentation, evaluation results, and reproducibility materials should be funded through academic, nonprofit, and national laboratory programs, with full release as a condition of support. Give priority to research where a model's output is a claim others must be able to reproduce, and to missions where an agency must know how a model was trained.

Turn Leadership into Adoption

Technological leadership matters only if American AI becomes the default choice for governments, companies, and developers. At home, that means reaching beyond the federal enterprise to the schools, hospitals, laboratories, statehouses, and small businesses where most of American life and work actually happens. Federal agencies should set the example as early customers and demanding users whose deployments validate new technology and reduce the uncertainty for those that follow.

1. Justified confidence that systems will work for their intended purpose

A major obstacle to deploying open models in government is justified confidence: what is the evidence that a model works reliably in a particular operating environment, performs measurably against alternatives, with performance tracked over time and failure modes understood well enough to set a tolerance. Safety testing and evaluation asks whether a model could cause harm; justified confidence asks whether it can do the

job under defined conditions. Institutions need both, but the second remains underdeveloped.

Part of the justified confidence gap is a perception that an open model cannot be secured. It can. Open models can carry the same protections as closed ones: guardrails, meaning built-in limits on what a model will do; classifiers, meaning separate filters that screen requests and responses for prohibited content before they reach the model or the user; and post-training safeguards, meaning the additional training a developer applies after a model is built to teach it to refuse harmful requests. An organization running an open model can inspect all three and verify that they are present, which it cannot do with a closed model. An open model can be certified for use in government environments, and it can run on an agency's own hardware, in a walled-off environment, so that sensitive data never leaves the building.

The National Security Commission on Artificial Intelligence identified the need for justified confidence five years ago. If AI systems routinely do not work as designed, or are unpredictable in ways that carry significant consequences, then “leaders will not adopt them, operators will not use them, Congress will not fund them, and the American people will not support them.”²¹

The consequence is predictable. Agencies default to the frontier models they already know, not because open models perform materially worse—they lag the frontier by only roughly four months—but because there is no mechanism to demonstrate fitness for a specific use. An official who cannot show that an open model meets a requirement will choose the one with the recognized name.

2. Support the open model ecosystem in public services and local institutions

Highways, the electric grid, and the internet became strategic assets when they reached every community, not when they were invented. An open American model ecosystem can play the same role. Cost, privacy, and institutional control make open models especially valuable to the institutions that shape American life and would not negotiate a frontier-lab enterprise contract.

²¹ [Final Report](#), National Security Commission on Artificial Intelligence (2021).

State and local institutions deliver much of American government while small and midsize firms account for a large proportion of employment. Most of these institutions and organizations cannot afford a dedicated AI team and none should have to navigate adoption alone. A school district should be able to run instructional systems locally without sending student records to a third party. A community hospital should be able to fine-tune a model on clinical workflows while keeping patient data inside its own walls. A county government, a rural utility, and a community bank should each be able to deploy on their own hardware, under their own control, with their own data.

These institutions face a decision problem before a technical one. A chatbot is easy to understand: you type a question, you get an answer. The best uses of an open model, and the gains for an organization from running its own models, are far less obvious. Part of what an adoption program has to deliver is a simple to understand explanation of which model best fits which use case.

The instinct to build regionally-delivered technical expertise, modeled on agricultural extension or the Manufacturing Extension Partnership, is right. But these institutions should not stop at advice. A program that teaches organizations how to use open models is worth far less than one that aggregates knowledge (e.g., what has already worked) and carries it between them. A written case study is worth less than a working package (i.e., the model, its configuration, the evaluation run against it, and the deployment steps) that the next institution can run instead of rebuilding. Local government is the natural network for that: a county that solves a problem has no commercial reason to withhold the solution from the next county. An extension capability built on that premise generates the record of results that adoption elsewhere depends on.

3. Make the federal government an early customer and validator

Open models let federal agencies run a model on their own infrastructure, adapt how it behaves to their mission, and keep sensitive data in-house. When the developer also documents what the model was trained on and how, the agency gets provenance, a verifiable record of where the model came from, on which accountability and independent review depends.

agencies have offices that publish what has worked and offer help; for open models, there is no equivalent, so agencies considering a deployment have to find each other. Nor does anyone measure whether open models are saving money or improving results, even though the tools to measure that already exist. Agencies keep portfolio scorecards that rate programs on cost, schedule, and performance, and government-wide IT is already measured under FITARA.²² Neither says anything about open models. Adding open-model deployments to those instruments (whether cost per task fell, whether the work got done faster, whether the mission outcome improved) adapts what exists to produce the evidence an agency needs to expand a deployment or end it, and gives the agencies that do well a way to claim credit.

Recommendations

1. Fund justified-confidence tooling as a capability distinct from safety evaluation.

Institutions need a way to establish that a model works for their purpose and keeps working. NSCAI identified this requirement in 2021; the tooling was never built.

2. Establish and fund an extension capability built around aggregating and sharing working deployments rather than only delivering advice. The product is a transferable record of results and, where it can be assembled, a working package the next institution can run.

3. Measure open-model deployments on already-existing agency scorecards. Add cost per task, time to complete, and mission outcome for open-model deployments to existing portfolio scorecards and FITARA reporting, and publish the results so the record is available to the next agency and to the extension network.

Govern American AI Without Making It Inaccessible

The diffusion race cannot be won by ignoring risk and harm. It can also be lost if governance treats openness itself as the risk. American open weight developers carry costs that foreign open-weight competitors may not. Meeting regulatory and

²² GAO-22-106105, [Information Technology and Cybersecurity: Using Scorecards to Monitor Agencies' Implementation of Statutory Requirements](#), U.S. Government Accountability Office (2022).

customer requirements, building and testing guardrails, and evaluating a model for safety all take money and time, and each adds to what a model costs and how long it takes to release. A competitor facing fewer of these requirements can release sooner and charge less, and a cheaper model that a user can download is easier for a developer to adopt. Governance investments build trust and long-term advantage, but poorly designed rules widen the price and convenience gap that drives adoption away from U.S. technology.

1. Openness is not a risk category

Open models differ in capability, release structure, intended use, access controls, built-in safeguards, and downstream impact. Two open-weight models may differ from each other far more than either differs from a closed system. A rule keyed to whether weights are public does not address risk. It is simply a rule against open models.

Open models are routinely discussed as though they are all Chinese models, and Chinese models as though they were all unsafe. The first association is easy to form, because the most capable open models available today are largely Chinese. The second might also be true, but the conclusion—all open models are unsafe—does not follow. A model whose weights can be inspected can be examined for exactly the properties a closed model asks its users to accept on trust. The assumption that open means unsafe has been repeated far more often than it has been tested. No independent body of research establishes what risks open models actually carry: how they fail in deployment, how well their safeguards hold up, or what capability a given release adds that was not already available. Rules should not be written on the strength of an intuition.

Some risks are real and should be stated plainly. Once released, weights cannot be recalled. Safeguards trained into a model can be removed by users with modest compute.²³ Provenance and training data are not always independently verifiable. Self-hosting shifts security burdens to the deployer. The useful question about any release is not whether it is open but what marginal risk it introduces: whether it makes available a capability that was not already available, from open or closed systems. A

²³ Xiangyu Qi, et al., [Fine-Tuning Aligned Language Models Compresses Safety Even When Users Do Not Intend To!](#), ICLR Conference Paper (2024).

release that changes nothing about what a determined actor can already do has not added risk, whatever its license says.

2. Govern high-consequence use cases through the regulators who own them

SCSP has held a consistent position since 2022. Chapter 3 of *Mid-Decade Challenges to National Competitiveness* sets out four principles: govern AI use cases and outcomes by sector; empower and modernize existing regulators; focus governance on high-consequence use cases; and strengthen non-regulatory governance.²⁴

The Framework for Identifying Highly Consequential AI Use Cases, developed in collaboration with the Johns Hopkins University Applied Physics Laboratory, is a template regulators can adapt to their own sectors.²⁵ It does not list technologies. It asks a regulator to look at a particular use of AI (e.g., a model approving loans, reading medical scans, screening job applicants) and to weigh the harms and benefits that use could produce, how likely each is, and how many people it would reach. A use whose failure would be severe, widespread, or hard to reverse is highly consequential and warrants the regulator's attention; most uses are not and do not. The regulator applies it when a new application appears in its sector, or when an existing one turns out to matter more than expected, and the result is a public registry of uses that require focus. What counts as consequential in aviation is not necessarily consequential in banking or education, so each regulator has to build its own version. But the method carries across sectors: consequence depends on where a model is deployed and what it is deployed to do.

“Open-weight” is exactly the kind of technology list item the Framework rejects. The regulators who already own each domain (e.g., the Food and Drug Administration, the Securities and Exchange Commission, the Federal Aviation Administration, banking and education authorities) are positioned to judge deployment context, and they already govern outcomes in their sectors regardless of the technology producing them. However, they cannot do it without technical capacity. A framework that relies

²⁴ [Mid-Decade Challenges to National Competitiveness](#), Special Competitive Studies Project (2022).

²⁵ [Framework for Identifying Highly Consequential AI Use Cases](#), Special Competitive Studies Project, Johns Hopkins University Applied Physics Laboratory (2023).

on existing regulators requires those regulators to employ people who understand the systems they are judging.

3. Risk and harm reporting is the nearest thing to a ready answer

When AI risks materialize, there is no standardized way to report the risk, the harm, the cause, or the solution. A 2026 study of AI vendors' vulnerability disclosure policies found wide variation in quality and scope: many vendors lack clearly defined disclosure channels, some restrict reporting to narrow categories of conventional security vulnerabilities, and some pathways are invite-only or exclude AI-specific flaws entirely.²⁶

The function that is missing is not complicated: somewhere for a disclosure to be sent, personnel able to assess them, and a mechanism by which the rest of the ecosystem learns about it, including potential mitigations and remedies. Legislation describing a voluntary reporting capability at NIST, in consultation with CISA, with common definitions, classification criteria, technical standards, and a national database, has been reported out of committee unanimously.²⁷ Whatever becomes of that bill, the function it describes is the gap.

Design is not the whole of it. A reporting body that is understaffed will not produce the adjudication the system needs, and an evaluation that is run, filed, and forgotten produces no resilience. It must be appropriately staffed and resourced.

Recommendations

1a. Avoid any policy approach that treats openness as a proxy for risk. Hold open and closed models to the same evaluation and provenance standards when used for similarly sensitive tasks. Requirements should be proportionate to capability and deployment context. Provide shared testing resources and technical support so smaller developers can meet these standards without being priced out.

1b. Fund independent, published research on the actual, not suspected, risks of open models. Research should be supported through NIST, the national laboratories, or NSF

²⁶ [FLARE-AI: Flaw Reporting for AI](#), Arvix (2026).

²⁷ H.R. 9333, [AI Flaw Reporting and Security Enhancement Act](#), 119th Congress (2026).

into how open models fail in deployment, how their safeguards hold up under attack, and what marginal capability a release adds beyond what is already available. Rules keyed to openness should rest on that evidence or not be written.

2. Direct sector regulators to develop frameworks for identifying highly consequential AI use cases in their own domains, using SCSP's Framework as a model, in place of rules keyed to release category, and fund the technical expertise inside those regulators that applying them requires.

3. Establish a standardized capability for reporting, adjudicating, and disseminating AI risks and harms, with common definitions, a national database, and the staff to run it.

The Metric That Matters: Who Is the World Building On?

The United States has the assets: frontier laboratories, deep capital, research universities, cloud infrastructure, national laboratories, and a network of allies no competitor can match. What it lacks is a strategy that turns those assets into defaults—the models developers reach for, the tools they learn, and the standards embedded in what gets built next.

China is running that playbook now, and defaults harden. Once developers around the world have built on a model, learned its tooling, and wired it into their infrastructure, displacing it becomes expensive. America does not need to copy that strategy. It needs a better one, and it needs to run it while the choice is still open.

The three lines of effort in this memo form an open-model diffusion strategy. Fund the open model ecosystem, through a Genesis Mission resourced to its ambition and a fully open model program in the fields where reproducibility matters. Buy it and measure it, so that federal demand reaches the companies that build American open models and the results travel to the next agency and the next county. And govern the use, not whether a model is open, through the regulators who already own each sector and the reporting infrastructure they need to do it.

This competition will not be decided by who builds the best model. It will be decided by whose models the world builds on, and by which society puts them to work first. That can and should be the United States.

Acknowledgements

On September 21, 2026, the Special Competitive Studies Project hosted a half-day workshop, *The Open Frontier: Winning the AI Diffusion Race*, under the Chatham House Rule. The participants included companies, frontier model representation, government officials, and academia. SCSP thanks the participants for their contributions. SCSP staff prepared this report and recommendations, drawing on many, though not all, of the insights from that discussion.

Ylli Bajraktari

President

Nyah Stewart

Director

Rama Elluru

Senior Advisor

Ananmay Agarwal

Director

Jenilee Keefe Singer

Senior Director

Chase Whitehouse

Associate Director

David Danks

SCSP Advisor