

Trump and the Tech Titans: Who Will Control the Age of Super Intelligence?



President Donald J. Trump brings together some of the world's most influential technology leaders as Washington confronts the opportunities and risks of the emerging age of super intelligence.

President Donald J. Trump brought together some of the most powerful figures in global technology to confront a question that could define this century: how far and how fast should artificial intelligence advance? With Elon Musk, Mark Zuckerberg, Jensen Huang, Sundar Pichai, Satya Nadella, Jeff Bezos, Dario Amodei, Greg Brockman, Lisa Su, and other technology leaders at the table, the White House meeting produced a voluntary framework for governing frontier AI. It also revealed something larger. Artificial intelligence has moved beyond Silicon Valley. It is now a matter of corporate governance, national security, industrial policy, and geopolitical power.

When President Donald J. Trump gathered some of the most powerful figures in global technology around a White House table, the occasion was about considerably more than another meeting between Washington and Silicon Valley.

Around the President sat executives whose companies control many of the essential building blocks of the emerging artificial intelligence economy, from frontier models and semiconductor chips to cloud computing, cybersecurity, enterprise software, social platforms, and the enormous data centers required to power them.

The significance of the gathering lay in the extraordinary concentration of technological and economic power represented in one room. These were not simply software executives discussing a new generation of consumer products. Their companies are collectively determining how trillions of dollars of future investment may be allocated, where data centers will be built, which countries will possess the most advanced computing infrastructure, how businesses will reorganize themselves around intelligent machines, and, ultimately, how much autonomy society should permit increasingly capable computer systems to exercise.

Trump approached that question from a position that has become central to his second administration. He regards advanced artificial intelligence as both an economic revolution and a strategic contest in which the United States cannot afford to surrender technological leadership. After the meeting, he said some people believed the technology could become bigger than the Industrial Revolution and repeatedly stressed his determination to maintain America's advantage. He also made clear that his preferred model was not an immediate expansion of federal regulation but rather the placement of substantial responsibility on the companies developing the technology themselves.

That philosophy produced the most important outcome of the meeting, the White House Accord on Super Intelligence and its "Joint Commitment on Frontier Responsibilities." Trump described the agreement as morally binding and compared its significance to a constitution. House Speaker Mike Johnson characterized it more precisely as a voluntary statement of principles and standards rather than a legally enforceable regulatory regime.

The distinction matters. Washington has not created a conventional licensing authority for frontier artificial intelligence. It has instead persuaded several of the companies closest to the technological frontier to accept a common architecture of corporate responsibility. The model is built around four layers of control,

progressing from the engineers and internal safety systems closest to the technology to independent external evaluation and, ultimately, the boardroom.

Trump's argument is that the people building these systems understand both their capabilities and their potential dangers better than almost anyone else. His preferred solution is therefore extensive self-policing backed by existing government authority rather than a regulatory structure that could slow development. He told reporters that there was strong support in the room for self-regulation and argued that ordinary federal law enforcement and government authority would continue to exist alongside it.



Elon Musk, Founder, xAI.

It is a distinctly American solution to one of the most complicated governance problems of the century. It seeks to preserve speed while building safeguards around it. Whether it proves sufficient will depend less on the words of the accord than on the seriousness with which companies implement them.



Mark Zuckerberg, CEO,
Meta.

Mark Zuckerberg emerged as one of the principal public explainers of what had been agreed. The Meta chief said the industry had developed principles centered on strong internal controls capable of detecting problems with advanced systems. Those controls would then be reinforced by successive levels of review, including internal risk assessment, outside auditors and evaluators, and independent consideration by company boards. Zuckerberg described the agreement as a starting point around which the wider industry could potentially converge.



Jensen Huang, founder and CEO of Nvidia, speaks in Stanford University's CS 153 course on April 30, 2026.

For Meta, the position is particularly significant because Zuckerberg has resisted the idea that leading AI laboratories should collectively stop or dramatically slow development. His broader argument has been that individual companies have a responsibility to develop models at a speed consistent with their ability to do so safely. The White House framework largely accommodates that philosophy. It does not require the industry to stop. It requires companies to demonstrate increasingly sophisticated systems of responsibility while continuing to build.




Sundar Pichai, CEO, Google.

Sundar Pichai placed the agreement in the context of corporate governance. He compared the proposed processes and controls with established financial controls inside major companies, effectively arguing that frontier artificial intelligence should become subject to similarly institutionalized disciplines. Google has described the accord as providing tangible mechanisms for responsible development while preserving the economic and scientific benefits of increasingly powerful systems.



Dario Amodei, CEO,
Anthropic.

That analogy may prove important. Financial controls did not eliminate financial risk. They created defined responsibilities, reporting structures, independent verification, and accountability extending to boards. The White House framework attempts to bring a comparable institutional structure to frontier computing.



Silicon wafer.

Jensen Huang represents another crucial part of the equation. NVIDIA is not principally competing with OpenAI, Anthropic, Google, or Meta by building the same kinds of consumer-facing frontier models. It supplies much of the computational infrastructure on which the AI revolution depends. Huang has consistently emphasized the transformative economic potential of accelerated computing and artificial intelligence. Following Trump's adoption of the term "Super Intelligence," Huang described modern data centers as becoming "superintelligence factories," an expression that captures the industrial scale on which AI is now developing.



Greg Brockman,
President and Co-
founder, OpenAI.

For corporate leaders, that description deserves attention. Artificial intelligence is moving beyond the economics of conventional software. Its expansion increasingly requires physical industrial infrastructure, including semiconductor fabrication, servers, networking systems, electricity generation, transmission capacity, cooling, land, and enormous capital expenditure. AI may be digital in output, but its foundations are becoming intensely physical.



Microsoft CEO Satya
Nadella

That was another major subject at the White House. Trump strongly defended the continued construction of data centers, while acknowledging the growing concern of communities facing increased electricity demand and other consequences of massive computing projects. He said technology companies had agreed to work more closely with local communities and suggested they could contribute to schools, teachers, energy provision, and other local needs.

The politics of AI infrastructure may consequently become as important as the technology itself. A frontier model cannot be trained merely through clever algorithms. It requires extraordinary quantities of computation. Computation requires chips. Chips require fabrication capacity and supply chains. Data centers require land, electricity, and water. Electricity requires generation and transmission infrastructure. Every stage creates a new intersection between technology companies, governments, and communities.

Elon Musk brought yet another perspective to the table. Few technology leaders have warned more publicly about the long-term risks of highly capable artificial intelligence, yet Musk is simultaneously one of the people building it through xAI. That apparent contradiction illustrates the central dilemma confronting the entire industry. Executives can simultaneously believe that advanced AI represents an extraordinary commercial and scientific opportunity and that sufficiently capable systems could pose unprecedented risks.

Musk's participation in the accord, therefore, matters. The agreement does not attempt to settle philosophical arguments about whether superintelligent machines could eventually exceed human control. Instead, it concentrates on operational safeguards that companies can implement now, including monitoring, independent evaluation, and board responsibility.



Jeff Bezos, Founder, Amazon.

Anthropic's Dario Amodei represents the more cautious end of the debate. He has argued publicly for slowing or pacing frontier development as capabilities become more powerful and potential dangers more serious. His presence and signature are therefore particularly noteworthy because the White House framework stops well short of the broader regulatory intervention that some safety advocates have proposed. The accord nevertheless incorporates several principles that cautious developers have long advocated, especially systematic testing, external evaluation, and governance extending beyond the engineers directly developing a model.

OpenAI occupied another important place in the gathering, although CEO Sam Altman was not at the luncheon. OpenAI was represented by its co-founder and President Greg Brockman, who signed the agreement. Altman separately said that OpenAI was increasingly willing to pace development when necessary and that there could be circumstances in which the company would choose not to train a model until it could make sufficiently strong safety arguments.

That position reveals how much the AI debate has changed. Only a few years ago, competition centered overwhelmingly on capability. Which model could write better, reason better, code better, or process more information? The next competitive frontier may increasingly require demonstrating that a powerful system can be controlled, audited, and trusted.

Satya Nadella attended the White House meeting but was not among the seven names appearing on the signed accord. The distinction is important. Microsoft is deeply embedded in the AI economy and has major relationships with frontier model developers, but attendance at the luncheon should not be taken as a signature on this document. The same applies to Jeff Bezos, Lisa Su, Alex Karp, and

several other major figures present at the gathering.

The final accord bears the names of President Donald J. Trump and six technology leaders: Sundar Pichai of Google, Dario Amodei of Anthropic, Mark Zuckerberg of Meta, Greg Brockman of OpenAI, Elon Musk of xAI, and Jensen Huang of Nvidia.



Lisa Su, CEO, AMD.

In substance, their Joint Commitment begins with a straightforward proposition. Companies developing advanced technology have a responsibility to build it safely and, in a manner, capable of earning the confidence of customers and the public. Frontier-model developers should therefore maintain sufficiently strong internal procedures to ensure their systems operate as intended and that problems are detected and corrected rapidly.

The first layer of the commitment concerns the technology itself. Each company undertakes to establish robust internal mechanisms to monitor the capabilities and alignment of frontier models during both training and deployment. Particular attention is directed towards cybersecurity, biological risks, chemical threats, and the possibility that models could penetrate or interact with computer systems in ways their developers did not intend.

This is perhaps the most technically significant section of the accord because it acknowledges that the principal danger may not simply be what a model tells a human user. Increasingly autonomous AI systems can use software tools, interact with networks, write and execute code, communicate with other systems, and undertake sequences of actions. Safety, therefore, has to move beyond controlling undesirable answers toward controlling undesirable behavior.

The second layer creates internal institutional accountability. Companies are

expected to establish or empower teams responsible for verifying that monitoring, detection, and control systems are functioning correctly and that identified problems are corrected. In other words, the people building a model cannot be the only people deciding whether their safeguards work.

The third layer moves outside the company's normal chain of command. Participating companies undertake to use independent external auditors or evaluators capable of assessing whether their controls and monitoring mechanisms are functioning as intended. This represents a potentially important shift in AI governance. External evaluation could eventually develop into an industry comparable in some respects to financial auditing, cybersecurity testing or certification, although the accord does not itself establish detailed universal standards for who qualifies as independent or precisely what must be disclosed publicly.

The fourth layer takes responsibility into the boardroom. Participating companies are expected to establish an independent committee of their boards to oversee the process, receive reports from internal teams and outside evaluators, and ensure that identified deficiencies are addressed.

That provision may ultimately prove the most consequential for corporate governance. AI safety would no longer be exclusively an engineering issue. It would become a board responsibility.

For chief executives and directors outside Silicon Valley, the implication is considerable. If AI becomes embedded in financial services, healthcare, manufacturing, logistics, defense, media, professional services, and government, boards may increasingly be expected to understand not only how their organizations use artificial intelligence but also how the associated risks are identified, independently tested, and escalated.

The commitment goes further. Participating companies intend to meet regularly to develop standards and best practices as the technology evolves. It also explicitly leaves open the possibility that the voluntary measures could eventually be incorporated into legislation or regulation. The companies nevertheless commit themselves to implementing the framework, whether or not the government ultimately makes it compulsory.

That sentence reveals the political compromise at the heart of the White House

gathering. Trump has not accepted the proposition that Washington should immediately construct a heavy regulatory apparatus around artificial intelligence. The companies, however, have accepted that voluntary governance today does not necessarily exclude legislation tomorrow.

The President's approach rests on an economic calculation. Artificial intelligence is no longer viewed merely as another technology industry. It is becoming a foundation technology capable of affecting productivity across almost every sector. Washington is simultaneously considering safety, national security, electricity supply, semiconductor production, data-center construction, education, employment, and competition with China.

Trump has repeatedly framed that competition in national terms. His administration wants the United States to retain leadership in frontier models, semiconductor technology, computing capacity, and the physical infrastructure supporting them. From that perspective, excessive regulation carries its own risk. A safety regime that substantially slows American development while competitors continue to advance could have economic and national-security consequences.

Critics of self-regulation raise the opposite concern. Companies engaged in an extraordinarily expensive technological race have powerful incentives to move quickly. Frontier AI laboratories compete for talent, investment, customers, and technological prestige. The danger is that commercial pressure could overwhelm voluntary restraint precisely when restraint becomes most necessary.

The White House Accord attempts to bridge those positions without resolving the fundamental disagreement between them. It allows development to continue while trying to institutionalize checks around that development.

There is also a deeper transformation taking place. For most of the digital era, governments regulated technology after markets had already developed. Social media became globally dominant before governments fully understood its political and social implications. Digital advertising transformed media before regulators adapted. Smartphones rewrote commerce, communications, and privacy while legal systems struggled to catch up.

With frontier artificial intelligence, governments are attempting something different. They are discussing governance while the technology itself is still rapidly evolving.

That creates an unusual relationship between political and corporate power. Governments possess legal authority but frequently lack the technical knowledge and operational visibility of the companies building frontier systems. Technology companies possess the expertise but cannot independently claim democratic legitimacy to determine the rules governing technologies that may affect entire societies.

The meeting represented an attempt to construct a bridge between those two centers of power.

For CEOs watching from outside the technology industry, perhaps the most important message is that AI has entered the boardroom permanently.

The first phase of the generative AI revolution was about experimentation. Companies asked what chatbots could do, whether employees should use them, and how much productivity could be gained.

The next phase is about integration. AI systems are becoming connected to corporate databases, customer information, software development environments, financial systems, and physical operations. Agents are increasingly capable of taking actions rather than simply generating text.

The phase now emerging is governance.

Who is responsible when an autonomous system makes an unexpected decision? Who verifies that a company's safeguards work? What information reaches the board? When should an external evaluator be brought in? How should cybersecurity risks created by increasingly capable models be measured? When does competitive speed become unacceptable operational risk?

These are no longer questions exclusively for AI laboratories.

The Trump meeting may therefore be remembered less for the terminology of "Super Intelligence" than for the governance structure that accompanied it. Internal controls. Independent internal oversight. External evaluation. Board accountability.

Those four layers form a remarkably conventional corporate answer to an unconventional technological problem.

Whether they will be enough is impossible to know. The accord is voluntary. The

frontier is moving extraordinarily quickly. There is no guarantee that every important AI developer worldwide will follow the same rules. There are unresolved questions about auditor independence, disclosure, common technical standards, and what happens when an evaluator concludes that a model is unsafe but commercial pressure favors release.

Yet the agreement marks a notable moment in the relationship between Washington and Silicon Valley. Some of the most influential figures in technology have lent their names to the proposition that building increasingly intelligent machines entails responsibilities that extend beyond ordinary product development.

Trump, meanwhile, has placed his administration firmly behind continued American technological acceleration while asking the companies driving it to construct their own layers of restraint.

It is an unusual bargain.

Washington will try not to suffocate the frontier. Silicon Valley will be expected to demonstrate that the frontier can be managed.

For Musk, Zuckerberg, Pichai, Huang, Amodei, and Brockman, the signatures represent more than corporate symbolism. They attach some of the best-known names in global technology to a common proposition: companies creating the most powerful computational systems must also accept responsibility for controlling them.

For Trump, the calculation is larger. He sees the technology as a foundation of future American economic and strategic power and has little interest in allowing fear alone to determine the pace of development.

And for the corporate world, the message from the White House table is difficult to ignore.

Artificial intelligence is moving beyond the technology department.

It is becoming infrastructure, capital expenditure, national strategy, corporate governance, and geopolitical power.

The race to build intelligence has already begun.

The race to govern it has now begun as well.

Joint Commitment on Frontier Responsibilities

The signatories state that companies developing frontier intelligent systems have a responsibility to develop them safely and in ways that sustain public and customer trust. Developers should maintain effective internal processes to ensure their systems behave as intended and that emerging problems are quickly identified and corrected.

They commit to four principal levels of protection.

First, developers should maintain strong internal controls throughout model training and deployment. These controls should monitor capabilities and alignment, with particular attention to cybersecurity, biological and chemical dangers, and unintended access to computer or technical systems.

Second, each company should empower a separate internal function to verify that those controls, monitoring systems, and detection mechanisms are operating correctly and to ensure that problems are remedied.

Third, each company should engage an independent external auditor or evaluator to assess whether those safeguards and monitoring arrangements are functioning as intended.

Fourth, each company should establish an independent board committee responsible for oversight. That committee should receive information from the teams administering the controls, as well as from internal and external auditors or evaluators, and ensure that identified weaknesses are addressed.

Collectively, the four layers are intended to increase confidence among companies, customers, and the public that advanced systems are operating as intended.

Participating companies also commit to meeting regularly to develop evolving standards and best practices as frontier systems become more capable.

Finally, the commitment recognizes that governments may eventually decide to place some of these practices into law or regulation. The participating companies state that, irrespective of whether such legal requirements emerge, they regard the controls and audit structure as important to the safe development of advanced

systems and commit themselves to implementing them.