CoreWeave Statement for The Record
Winning the AI Race: Strengthening U.S. Capabilities in Computing and Innovation
Senate Committee on Commerce, Science and Transportation
May 8, 2025

Introduction

Chairman Cruz, Ranking Member Cantwell, and distinguished Members of the Committee, I want to thank you for the opportunity to provide testimony today on how we can ensure the United States remains the global leader in Artificial Intelligence (AI) innovation.

Today we're on the verge of the next great revolution of AI: a technology dramatically reshaping industries, innovation, and productivity at a massive scale through systems of unprecedented complexity. Millions of hours of training, billions of inference queries, trillions of model parameters, and continuous dynamic scaling are all driving an insatiable hunger for compute and energy that borders on exponential.

At CoreWeave, we are not only deploying the infrastructure that hosts some of the most massive and powerful AI models in existence, but also doing it at scale.

I appreciate the opportunity to describe CoreWeave's journey and our role in this historic race.

CoreWeave's Role

CoreWeave powers AI innovations by bridging the gap between AI ambition and execution. CoreWeave provides the cloud platform, purpose-built for AI, that delivers the speed, performance, and expertise needed to unleash AI's full potential.

To train their next generation of AI models, AI labs require compute resources. And to maintain their leadership position at the forefront of AI innovation, researchers demand the latest and most-performant computing infrastructure. They leverage these compute resources to run the trillions of operations that operate algorithms and process data to train the next generation of models.

The "Scaling Law"[1] has demonstrated that increasing the compute deployed against models translates to better performance. The relationship is exponential. Orders of

[1] Scaling laws describe how the performance of AI systems improve as the size of the training data, model parameters, or computational resources increases.

magnitude of increased compute are required to unlock incremental gains in model performance. For example, in 2018, one of the most popular and advanced generative AI models required a certain amount of compute to train. Just seven years later, in 2025, the amount of compute needed for the latest frontier models grew by about 100,000 times — an extraordinary increase.[2] Performance, usefulness, and real-world adoption has increased dramatically. This demonstrates that the limiting boundary of AI development is high-performance compute, delivered at scale and operated at peak efficiency.

To access the increased need for compute and to avoid being left behind, enterprises and labs are continuously demanding deployments of the latest generation chips at larger scales. In just the two short years between the most recent generations of Nvidia chips, training performance increased by 4x from Hopper to Grace Blackwell.[3]

CoreWeave enables AI labs, platforms, and enterprises to move the boundaries of compute forward. Our end-to-end cloud platform is purpose-built for the scale, performance, and expertise needed to power AI innovation and meet the demands of accelerated computing. We construct and power data centers, and provision high-performance computing infrastructure which enables enterprise and labs to access these essential resources. Our expertise is in managing a complex and fragile ecosystem across supply chain, energy, financing, and technology partnerships to build, optimize, and deploy our platform at scale.

As of December 2024, we operated in 32 data centers in the U.S. and Europe, deploying more than 250,000 GPUs utilizing 360MW of active power. We have now contracted for 1.3GW of power. Adequate, reliable supplies of power are essential to drive this revolution and for the U.S. to win this race.

This year, CoreWeave was the sole cloud to be ranked Platinum and as the #1 leader in AI cloud performance, as attributed by SemiAnalysis's Platinum ClusterMAX™ Rating.[4] And we have established a track record as among the first to market with the latest generations of hardware, such as Nvidia's most recent GB200 NVL72 chip, which leading labs, such as IBM, Mistral AI, and Cohere, are already using to improve and accelerate their training jobs.[5] We also support OpenAI-- our strategic deal of nearly

[2] Epoch AI, "Training Compute of Frontier AI Models Grows by 4-5x per Year," May 28, 2024.
[3] Nvidia, "NVIDIA GB200 NVL72: Powering the new era of computing."
[4] CoreWeave, "CoreWeave Ranks as #1 AI Cloud, Backed by SemiAnalysis' Platinum ClusterMAX™ Rating," April 10, 2025.
[5] Nvidia, "Thousands of NVIDIA Grace Blackwell GPUs Now Live at CoreWeave, Propelling Development for AI Pioneers," April 15, 2025.

$12B provides compute capacity for training and delivering its latest models at scale to its hundreds of millions of users around the world.[6]

CoreWeave is purpose-built for the demands of accelerated computing. We deliver this infrastructure with cutting-edge performance and scale, and provide the expertise with the infrastructure that AI needs today and in the future. And as a result, our customers are able to train earlier, build quicker, and get to market faster, which is critical for the U.S. to maintain its lead in AI. CoreWeave is the engine that will propel the U.S. forward in the AI race.

CoreWeave's History

CoreWeave began as many start-ups and great entrepreneurial companies do – in a garage with an idea, which was to try leveraging GPUs for crypto mining. CoreWeave's founders purchased their first GPU in 2016, which turned into hundreds, then thousands, and now hundreds of thousands. Over the course of the next few years, CoreWeave began looking for opportunities to use its fleet of GPUs for other high-performing use cases beyond crypto mining, such as visual effects (VFX), and then AI.

- In 2020, CoreWeave launched as the "world's first specialized cloud."
- In 2021, CoreWeave operated the largest Nvidia A40 fleet in North America.
- By 2022, the world began to realize that more compute was required to scale AI model training. CoreWeave began to scale even more rapidly.
- By 2023, the company had three data centers running more than 17,000 GPUs with approximately 10 megawatts (MWs) of active power.
- By 2024, CoreWeave had ten data centers running more than 53,000 GPUs with more than 70MW of active power.
- And, one year later, at the end of 2024, we had 32 data centers running more than 250,000 GPUs with approximately 360MW of active power.

CoreWeave became a publicly traded company on March 28, 2025.

This progress occurred in five short years. And that is the speed which is required to drive this technological revolution. Most recently, CoreWeave completed its acquisition of Weights & Biases, a leading AI developer platform. Our vision is that CoreWeave + Weights & Biases will deliver the leading AI Cloud Platform—purpose-built to develop, deploy, and iterate AI faster. Together, we will enable faster, more efficient AI

[6] CoreWeave, "CoreWeave Announces Agreement with OpenAI to Deliver AI Infrastructure," March 10, 2025.

development, empower AI developers to innovate quickly, provide seamless integrations for AI development, and support the world's most advanced AI innovators to unleash AI's full potential.

The Need for AI Infrastructure and Re-platforming

AI requires a fundamentally different computing infrastructure from the existing one. Training state-of-the-art models and running inference at scale requires trillions of simultaneous calculations across billions of parameters. To fulfill this requirement, high-performing compute infrastructure necessitates a more concentrated power footprint, increased cooling needs, the latest chips, high-throughput networking, accelerated storage, and more.

In contrast, the generalized clouds that serve the world today were not built to serve the specific requirements of AI. These cloud platforms were built for day-to-day web hosting, database management, and running SaaS applications – workloads that rely on simple, fixed-logic calculations and lightweight processing.

Operating compute at scale and at the intensity of AI is highly complex. There are significant inefficiencies associated with operating AI workloads ranging from hardware failures to scheduling optimal usage. A single 32,000 GPU cluster may require the deployment of approximately 600 miles of fiber cables and approximately 80,000 fiber connections, along with highly-specialized heat management capabilities to support high-power density. Each of these variables increases the number and complexity of possible failure points. When a cluster suffers a component failure (GPU, network, memory, cable, cooling, etc), it can adversely impact the entire cluster by reducing training performance, or even causing the entire project to fail.

The difficulty in managing large clusters leads to what we call the "AI Efficiency gap," which we evaluate based on Model FLOPs Utilization (MFU). This is a measure of the observed throughput compared to system maximum if the system were operating at peak capacity. Typically, the complexity of managing AI infrastructure means that a majority of the compute capacity in GPUs can be lost to system inefficiencies, with empirical evidence suggesting observed levels of performance in the 35% to 45% range.

As a result, the world is undergoing a "re-platforming" from traditional generalized cloud computing infrastructure to AI cloud computing infrastructure. And to achieve this, cloud platforms are being fundamentally reimagined, with every layer of the technology stack

being specifically optimized for AI workloads. This is the purpose-built computing infrastructure needed to support the scale and complexity of AI workloads.

CoreWeave's Cloud Platform

We have built our platform for the new requirements of AI cloud computing infrastructure.

CoreWeave's cloud platform is an integrated solution that is purpose-built for running AI workloads such as model training and inference at superior performance and efficiency. It includes infrastructure services, managed software services, and application software services, all of which are augmented by our proprietary Mission Control and observability software. This proprietary software enables the provisioning of infrastructure, the orchestration of workloads, and the monitoring of our customers' training and inference environments to ensure high availability and minimize downtime.

To unlock the full potential of AI infrastructure, CoreWeave helps to bridge the MFU "efficiency gap" between the observed 35–45% and the theoretical 100%, driving as much as 20% higher performance than public benchmarks.[7] To achieve this, performance optimizations are built into every layer of the platform to enhance distributed training throughput. And our ability to close this gap significantly enhances performance, improves model quality, accelerates development timelines, and reduces overall AI model costs.

What does this mean for the U.S.? As we improve our efficiency, and close this gap, the United States will maintain its edge in the global AI race, stimulating economic activity, and enhancing national security while improving the provision of essential services for all. This is what the race is all about.

U.S. Global Leadership in AI

We stand at a critical inflection point in the global AI race, representing a pivotal moment that will influence economic prosperity, national security, technological standards and how we provide essential services to all Americans. AI represents the next major evolution of technology with the potential to transform society. This is America's AI moment, and a strategic opportunity America cannot afford to miss.

[7] CoreWeave, "CoreWeave leads the Charge in AI Infrastructure Efficiency, with up to 20% Higher GPU Cluster Performance than Alternative Solutions," March 19, 2025.

Economic Prosperity: AI is projected to generate a cumulative global economic impact of $20 trillion, representing 3.5% of global GDP, by 2030.[8] The country that leads this transformation will capture a disproportionate share of this new economic frontier. If America maintains global leadership in AI, the productivity gains, new products, high-value jobs, and breakthroughs across industries from healthcare to manufacturing created by AI will help drive prosperity across the American economy benefitting all people.

National Security: As advanced AI capabilities become essential to modern defense including improvements to weapons systems and battlefield capabilities, intelligence, and cybersecurity systems, maintaining America's technological edge becomes inseparable from our national security. Falling behind is not an option when other countries are rapidly advancing their own AI capabilities with explicit aims to challenge American global economic and military leadership.

Shaping the Future of AI: The country that leads AI development will shape how this technology evolves globally. The standards, protocols, and ethical frameworks that will govern AI will reflect the values of whichever country wins this race.

The foundational AI infrastructure being built today will help determine where AI development occurs. Success in the global AI race will increasingly depend on purpose-built AI computing infrastructure, not just general-purpose systems deployed at scale. Nations that successfully "re-platform" gain compounding advantages in model capabilities and development speed.

CoreWeave is at the forefront of developing the purpose-built infrastructure that powers America's AI capabilities. Leading companies and AI labs such as IBM, Mistral, and Cohere rely on CoreWeave's infrastructure. Our success supports broader national objectives by ensuring the U.S. maintains the world's most advanced computing infrastructure which is required to drive AI.

Factors Critical to Continued U.S. Leadership in AI Infrastructure

America's leadership position in AI depends in part on maintaining its edge in the underlying infrastructure that drives it. Based on CoreWeave's experience building and operating AI computing infrastructure, I would like to highlight several critical areas that will determine if our nation maintains its leadership position. Many of these areas focus on the critical elements of policy which will impact how this sector evolves.

[8] IDC, "The Global Impact of Artificial Intelligence on the Economy and Jobs," September 2024.

Strategic Investment Stability

AI infrastructure investment requires a significant level of coordination across multiple industry and government stakeholders due to the scale and timeline of these projects, representing substantial capital commitments with years-long development and operational horizons. CoreWeave benefits from robust collaborations with leading chipmakers, original equipment manufacturers ("OEMs"), and software providers to supply us with infrastructure components and other products. The highly specialized infrastructure that is required to unlock the potential of AI is immensely challenging to build and operate, especially at scale. This requires: (i) tens of thousands of GPUs, (ii) thousands of miles of high-speed networking cables, (iii) hundreds of thousands of interconnects coming together to create "superclusters" for training and serving AI models, and (iv) hundreds of MWs of power and substantial amounts of storage.

To sustain U.S. leadership in AI, it is important for U.S. AI cloud computing companies to maintain access to a reliable supply chain necessary to access all of the components necessary to develop and run the cutting-edge AI infrastructure. Acquiring these necessary high-performance components to power AI workloads requires managing a complex global supply chain and maintaining robust supply chain relationships. Continued engagement with leading global suppliers and strategic partners is vital to ensuring the continued operation, expansion, and rapid deployment of U.S. AI infrastructure and to uphold U.S. competitiveness. Predictable policy is essential for this.

Significant private sector investment and development has helped the United States establish an early and important lead in AI infrastructure. The U.S. accounts for roughly 40 percent of the global market for data center capacity, with six of the top ten markets.[9]

The importance of AI and U.S. leadership is not lost on our competitors. Intensifying global competition for AI infrastructure demands that this initial lead must be carefully and actively maintained. The European Union launched its AI Continent Action Plan in April setting out ambitious goals to triple data center capacity across member states in the next five to seven years. The EU also announced a €20 billion investment into five gigafactories–massive high performance computing facilities equipped with approximately 100,000 state-of-the-art AI chips–and reforms related to permitting, energy issues and water usage.[10] China has made its ambitions regarding AI clear

[9] Cushman & Wakefield Research, "2024 Global Data Center Market Comparison."
[10] European Commission, "The AI Continent Action Plan," April 9, 2025, https://digital-strategy.ec.europa.eu/en/library/ai-continent-action-plan.

through coordinated national strategies and streamlined deployment timelines that can compress years into months in their effort to shrink America's current AI lead.

Countries around the world are aggressively pursuing coordinated AI strategies, implementing policies which subsidize infrastructure, and accelerate deployment timelines. In this high-stakes environment, the capacity of American companies to build AI infrastructure swiftly and with assurance will be a decisive factor in the AI race and determine whether the United States retains its leadership position.

Sustained American leadership in AI infrastructure faces potential headwinds from multiple sources of uncertainty. These include volatility in the global supply chain for critical components, such as advanced semiconductors and networking equipment which can disrupt deployment timelines. These fluctuations can lead to delays or unanticipated cost overruns, adversely affecting American companies' ability to rapidly scale AI capabilities.

Changes in regulation at both the federal and state levels introduce substantial uncertainty for leading businesses making investment decisions. Changes in export controls, energy policies, the potential need to add gigawatts of power to meet increased demand, and the emerging landscape of AI-specific regulations at different levels of government also affect the pace and scope of infrastructure deployment. The lack of regulatory clarity can deter investment and slow down the innovation cycle. Additionally, American companies will be affected by the rules and institutions being developed around the world both in individual nations and in important forums in which key competitors participate that will govern the use of AI. As the world's dominant economic player and technological leader, it is important for the U.S. government to drive the rules which shape the future playing field for American companies.

Finally, the potential for a fragmented regulatory framework, with differing requirements from state to state and potentially at the federal level, poses a unique challenge. For instance, inconsistent definitions of key terms which define various AI activities across jurisdictions could force companies to navigate a complex web of compliance regimes for fundamentally similar activities. These types of policies would require participants in the AI infrastructure to consider designing alternative products and strategies to do business in different jurisdictions. This regulatory patchwork would lead to increased costs, operational inefficiencies, and ultimately, a competitive disadvantage for American companies in the global AI race. These uncertainties disproportionately affect newer entrants like CoreWeave and other specialized providers, potentially stifling the very innovation that drives American leadership in this critical sector.

To ensure the United States remains at the forefront of AI, American companies must lead AI infrastructure development. This requires a coordinated policy strategy to mitigate key uncertainties, many of which this section touches upon, maintain appropriate oversight, and create a stable, predictable policy environment that fosters investment, continued growth and innovation.

Energy and Infrastructure Development: Powering AI Leadership

The race to build AI infrastructure is fundamentally tied to the nation's ability to continue to develop a new generation of data centers that drive innovation, to bring them online and ensure there is sufficient electricity to power them. This will be affected by the processes and permitting systems that are used to develop data centers and to develop adequate, reliable supplies of new power generation and the interconnection and transmission systems capable of delivering it at pace and scale.

The dual challenges of adding new power supplies and streamlining infrastructure development, are not merely logistical hurdles, but critical factors that will determine whether America can maintain its global AI leadership. Failure to address these challenges effectively risks ceding ground to international competitors, particularly China, who are aggressively pursuing their own AI ambitions.

Energy considerations are critical to the development and operation of AI infrastructure. After a prolonged period of relatively flat electricity consumption, according to analysis, the U.S. is now experiencing a significant and accelerating increase in power demand. This surge is driven by several concurrent trends, including the onshoring of new manufacturing facilities, the widespread electrification of transportation and heating, and the growth in data centers.

AI computation is energy-intensive. Training large language models, running complex simulations, and deploying AI applications all require significant amounts of power. Widespread AI adoption will further increase this demand, even as companies continue to innovate and improve efficiency. According to a report released by the U.S. Department of Energy, data centers consumed approximately 4.4% of total U.S. electricity in 2023. This figure is projected to rise in the coming years, potentially consuming between 6.7% and 12% of total U.S. electricity by 2028.[11] This projected increase underscores the urgent need for policymakers at all levels of government to put policies in place that will enable the development of new power supplies and the

[11] "2024 United States Data Center Energy Usage Report," Lawrence Berkeley National Laboratory, December 2024.

infrastructure to deliver it. Given that these projects can cost in the hundreds of millions of dollars and years to implement, there is no time to lose in getting started.

The implications for global AI leadership are clear and consequential. Regions that can provide abundant, reliable, and cost-effective energy will attract billions of dollars of AI infrastructure investment. Conversely, energy constraints, whether in the form of limited supply, unreliable delivery, policy uncertainty, or prohibitive costs, can and will push development and the associated investment elsewhere.

CoreWeave's site selection consideration for data centers illustrates these priorities:

- Availability of abundant, reliable power, and where available, non-emitting sources
- Competitive rates
- Diverse energy sources
- Pathways for capacity expansion
- Efficient permitting processes that provide timeline certainty

However, obtaining the necessary approvals to build both energy infrastructure and data centers is often a critical bottleneck. Every month of delay represents lost ground in a field where the pace of innovation is measured in weeks, not years.

In particular, it is challenging to develop energy projects. Securing the necessary approvals and permits can take a significant amount of time. Inefficiencies in the permitting process can significantly impact both energy availability and whether attractive sites for data center development can move forward. The variability in permitting timelines across jurisdictions and the potential for multiple, sequential review processes and litigation can increase the time required to develop a project leading to delay or potentially stopping projects.

There will be challenges in streamlining the permitting and regulatory processes required to develop the energy and data center infrastructure necessary for the U.S. to maintain its leadership in AI. Goals in streamlining these processes include:

- Maintaining and growing a balanced portfolio of generation powered by diverse energy sources that can meet increasing demand to ensure availability and reliability at reasonable costs
- Expanding and modernizing the nation's transmission systems
- Providing developers of data center capacity and associated infrastructure with predictable timelines and reduced wait times for feasibility studies, interconnections, and builds

- Streamlining permitting processes while maintaining appropriate oversight

CoreWeave understands that the processes put in place to achieve these important objectives need to consider the views of key players that will make these investments and the communities in which these facilities are located.

We hope efforts to streamline the permitting process to enable the addition of new sources of generation and the transmission infrastructure to deliver it will receive attention by this Congress.[12]

We recognize that this issue will not be resolved solely at the federal level. All levels of government have a role to play in addressing the challenges in adding the necessary infrastructure to meet energy requirements. A coordinated effort amongst federal, state, and local government, industry, and other affected parties is required to address these interrelated challenges which include creating efficient, transparent processes that allow infrastructure to be built at the pace and scale that technological advancement requires and for the U.S. to maintain its dominant position in AI.

Global Diffusion of the American AI Stack

Global market access is a pivotal factor in determining which nation will lead in the AI domain. Export controls and trade agreements can be designed to achieve multiple objectives: they can facilitate legitimate market access for American businesses while also mitigating potential national security risks. However, controls that are not calibrated can inadvertently bolster foreign competitors by incentivizing AI development and deployment outside of the U.S., and competitors filling the void left by U.S. firms. This could result in the loss of technological expertise and economic benefits.

To bolster American AI leadership, export controls and international agreements should consider:

- **Precision Targeting of National Security Risks:** Controls should be focused on technologies, entities and nations that pose genuine and demonstrable threats to national security, with clear and specific parameters.
- **Supporting American Technological Leadership:** Restrictions imposed on U.S. technologies and where they can be exported should consider negative impacts on the ability of U.S. companies to compete in global markets. This includes considering the potential for retaliatory measures from other nations and the risk of creating a 'chilling effect' on investment and innovation.

[12] CoreWeave is a member of the Data Center Coalition. DCC's AI Action Plan submission includes additional discussion related to permitting and energy infrastructure, available here.

- **Strategic Alignment with Allies:** Close coordination with like-minded international partners is essential to ensure the effectiveness of export controls and prevent the fragmentation of the global AI market. Aligning with allies can also foster the expansion of a collaborative, secure, and trustworthy AI ecosystem.

These considerations are crucial in shaping the future landscape of AI innovation. A well-calibrated approach will ensure that the next generation of AI development is anchored in the United States, leveraging American infrastructure and expertise. Conversely, controls could inadvertently have the effect of limiting opportunities for the export of U.S. technology and expertise, with adverse impacts. A strategy that carefully differentiates between markets, tailors export restrictions to mitigating specific risks, and fosters international cooperation, can effectively protect national security while simultaneously enhancing America's ability to lead and shape the global AI diffusion race.

Public-Private Partnerships Accelerating American Innovation

A unique American advantage in the AI race is our ability to forge effective partnerships between government, industry, academia and elements of civil society. These collaborations combine the agility, ingenuity, expertise and resources of the private sector with the long-term vision of the public sector and the basic and applied research capabilities of academic and research institutions. This approach helps foster an innovation ecosystem that is difficult for competitors to replicate.

CoreWeave is proud to be a founding partner of the New Jersey AI Hub, along with Microsoft, Princeton University, and the New Jersey Economic Development Authority. This AI Hub will focus on research and development efforts, applications of AI in several industry sectors, and AI workforce development and education.[13]

CoreWeave is deeply committed to supporting AI education and research, and public-private collaborative partnerships. Similar partnerships across the country could further accelerate America's AI capabilities, and we encourage policymakers to explore this model of collaboration. These types of partnerships also accelerate AI and data center workforce development in the US. Worker shortages in the data center industry are becoming commonplace as the skills gap widens. Data center employers have struggled to find enough trained workers. Half or 50% of surveyed data center managers in 2020 reported having difficulty finding skilled workers to fill positions, and

[13] Governor Phil Murphy, "Governor Murphy, Princeton University, Microsoft & CoreWeave Cut Ribbon on Major Artificial Intelligence Hub," March 27, 2025, https://www.nj.gov/governor/news/news/562025/approved/20250327a.shtml.

71 percent continued to report being concerned about finding qualified staff in 2023.[14] A skilled and trained workforce is vital for the stability and expansion of AI data centers – which rely on specialized data center technicians, network and electrical engineers, cybersecurity professionals, and project managers. CoreWeave supports efforts to develop a domestic workforce comprised of the skilled workers required to meet the growing AI demand and to accelerate AI innovation while creating the skilled good-paying jobs of the future.

Looking Forward: Ensuring Continued American Leadership in AI

The United States has built a remarkable lead in artificial intelligence through our unique combination of innovative and entrepreneurial private companies, world-class research institutions, a talented workforce, and a policy environment that fosters dynamic growth. This advantage is especially pronounced in AI infrastructure, where companies like CoreWeave have established global technological leadership in this critical layer of the AI stack.

However, the conditions that have enabled this position must be actively maintained while being flexible in order to adjust to new technological developments and political considerations. Countries around the world rightfully recognize the strategic importance of AI and are making coordinated efforts to build AI infrastructure. The decisions we make now will help determine whether America can maintain AI leadership in the years to come.

The current moment demands a thoughtful, transparent, and predictable approach that maintains our competitive edge, and seizes future opportunities while addressing legitimate concerns. As we consider policy options to address this dynamic sector, we should be attentive to how different approaches affect the entire AI ecosystem, from established players to new entrants, from model developers to infrastructure providers like CoreWeave.

In order to further America's lead in AI, we encourage the federal government to consider policies that:

- Foster a predictable investment environment through the implementation of nationally consistent regulatory frameworks for areas most critical to strengthen competitiveness and drive innovation.
- Ensure that there are adequate, reliable supplies of power at the lowest possible cost through policies which enhance the ability to add generation and transmission to power next-generation AI infrastructure. Careful reforms of

[14] The White House, "AI Talent Report," January 14, 2025

existing permitting and regulatory processes that enable affected parties to participate in the process are needed for this to occur.
- Maintain global competitiveness and strengthen U.S. industry through strategically calibrated export policies that protect national security while supporting the diffusion of the American AI stack.
- Strengthen public-private partnerships, like the New Jersey AI Hub, that accelerate innovation across research, industry, and government.

CoreWeave appreciates the Committee's leadership on these critical issues and we look forward to working with the Committee as it develops policies enabling the U.S. to maintain its leadership in the AI race.