



September 2025

Securing the Backbone of Artificial Intelligence: Protecting Data Centers

Seungmin (Helen) Lee

Future Security

Last edited on September 03, 2025 at 12:35 p.m. EDT

Acknowledgments

Special thanks to subject matter expert interviewees Austin Carson, Tim Fist, Allan Freedman, James Slaughter, and Ian Wallace; peer reviewers Michael Garcia and Adefoluke Shemsu; and advisor Peter Singer, program manager Bridget Chan, and others at New America's #ShareTheMicInCyber program for their support and contributions.

Editorial disclosure: The views expressed in this report are solely those of the author and do not reflect the views of New America, its staff, fellows, funders, or board of directors.

About the Author

Seungmin Helen Lee is a 2025 #ShareTheMicInCyber Fellow.

About New America

We are dedicated to renewing the promise of America by continuing the quest to realize our nation's highest ideals, honestly confronting the challenges caused by rapid technological and social change, and seizing the opportunities those changes create.

About Future Security

Future Security is a partnership between New America and Arizona State University. It reconceptualizes U.S. security policy towards a holistic engagement with current and future challenges including domestic terrorism, armed drones, climate change, pandemics, rising authoritarianism, and new and emerging technologies.

About #ShareTheMicInCyber Fellowship

The #ShareTheMicInCyber Fellowship invests, grows, and platforms emerging cybersecurity and technology leaders from traditionally underrepresented backgrounds in support of strengthening our nation's resilience to digital threats.

Contents

Executive Summary	5
Introduction	6
Methodology	8
The Evolution of Data Centers	9
A Brief History	9
Data Centers	9
AI Data Centers	12
Cyber Threats	14
Traditional Data Center Threats	16
AI Data Center Threats	19
Changing World Threats	22
A Framework for Cyber-Secure AI Data Centers	28
Technical Measures	28
Corporate Policy Measures	29
National Governance Measures	29
Conclusion	32
Appendix	33

Executive Summary

Global data center demand is predicted to triple by 2030, with nearly 70 percent of the demand being driven by artificial intelligence (AI) workloads. The surge in AI data center demand is fueled by the exponentially increasing use of AI at work and integration into daily lives as well as the recognition that the technology can influence geopolitics, reorder the global economy, drive scientific discovery, and transform human lives and society.

Two barriers to supporting this increased demand for AI data centers are the required energy and security. While the energy concern has been highlighted and publicized, the security gaps have been understated: AI data centers have mostly been considered as a part of traditional data centers and critical infrastructure, with no separate or focused effort for AI data center cybersecurity requirements.

However, AI data centers face an expanded set of threats. A successful cyberattack on an AI data center could enable threat actors to extract information about the AI model and weights, risking loss of sensitive training data as well as the integrity and confidentiality of the AI model. When AI models are exfiltrated, hackers can create vulnerabilities in AI models and bias outputs, as well as more easily and cheaply replicate AI models.

Thus, this report recommends a comprehensive framework for AI data center security that spans six layers of security and three types of approaches. The six layers of security are (1) hardware & compute, (2) network & storage, (3) model & data, (4) software & application, (5) physical access, and (6) geopolitical. Given the complexity, these six layers also require three approaches: technical, corporate policy, and national governance. Under this framework, the report's recommendations are as follows:

1. Bridge the gaps between the technical, corporate policy, and national governance approaches with a framework that maps the threats to AI data centers across the six layers of security.
2. Implement existing research and standards for technical requirements in an AI data center.
3. Require technical measures across the six layers of security in corporate policies.
4. Focus national governance measures on incentivizing operators to meet the technical and corporate policies needed for a cyber-secure AI data center.

Introduction

Despite the over 10,000 data centers already existing around the world, global data center demand is predicted to triple by 2030, with nearly 70 percent of the demand being driven by artificial intelligence (AI) workloads.¹ This surge in AI data center demand is fueled by the exponentially increasing use of AI at work and integration into daily lives as well as the AI race prompted by the recognition that AI technology can influence geopolitics, reorder the global economy, drive scientific discovery, and transform human lives and society.²

Beyond the predictions, 2025 started with announcements of investments into AI data centers: Project Stargate—supported by the Trump administration—pledged a \$500 billion investment for AI infrastructure; and Meta announced a \$60–65 billion pledge to AI and an AI data center.³ Globally, France announced that MGX, Bpifrance, Mistral AI, and Nvidia are collaborating on building Europe’s largest data center campus; EDGNEX Data Centres by DAMAC announced a \$2.3 billion investment into a 144-megawatt AI data center in Indonesia; and Chinese tech firms plan to build more than 30 AI data centers with 115,000 Nvidia chips in the Xinjiang region.⁴

Two barriers to supporting this increased demand for AI data centers are the required energy and security.⁵ Yet while the energy concern has been highlighted and publicized, the security gaps have been relatively neglected.

Data centers already make up 1 to 2 percent of global energy demand, and data center energy demand is projected to increase to 21 percent by 2030 due to AI.⁶ The International Energy Agency reported that in 2024 AI servers drove 15 percent of electricity demand from data centers, and it projected that data center electricity consumption will more than double by 2030, with AI as the most influential factor.⁷ Recognizing the energy demand, President Biden issued the Executive Order on Advancing United States Leadership in Artificial Intelligence Infrastructure in January 2025, which heavily focused on federal support for AI data center energy needs and clean energy.⁸ Unfortunately, President Trump revoked this executive order six months later.

On the security front, though, AI data centers have mostly been considered as a part of traditional data centers and critical infrastructure with no separate or focused effort for AI data center cybersecurity requirements. For example, the Biden administration issued its National Security Memorandum on Critical Infrastructure Security and Resilience (NSM-22), which designated data centers as a critical part of U.S. infrastructure, indicating the importance of defending the critical infrastructure sectors from cyber activity led by nation-states.⁹ Beyond the United States, China's 2016 Cybersecurity Law required data centers to have cybersecurity measures against domestic and foreign threats.¹⁰

However, AI data centers face an expanded set of threats. A successful cyberattack on an AI data center could enable threat actors to extract information about the AI model and weights, risking loss of sensitive training data as well as the integrity and confidentiality of the AI model.¹¹ When AI models are exfiltrated (stolen), hackers can create vulnerabilities in AI models, biased outputs, and similar AI models with fewer resources.¹²

Thus, this report recommends a comprehensive framework that identifies requirements for implementing sufficient cybersecurity measures in AI data centers. This report will first provide the methodology, followed by a brief history of traditional data centers and AI data centers. This background will help build a necessary foundation for the discussions of cyber threats to AI data centers. Finally, the report concludes with a framework and recommendations for a comprehensive cybersecurity approach.

Methodology

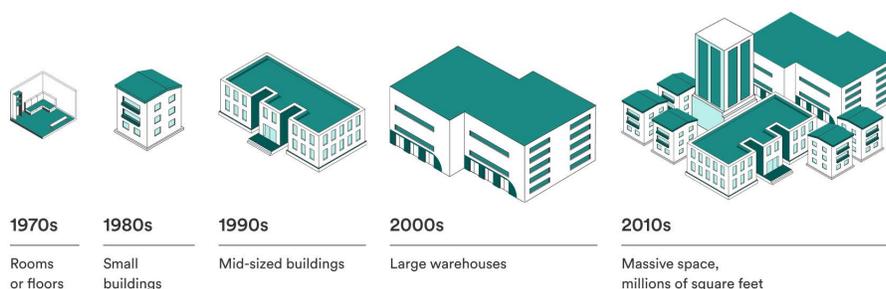
The findings in this report are based on a combination of open-source research, expert interviews, and analysis. The open-source research and expert interviews were conducted in parallel. Technical information was selected from technical reports, cyber and technology companies' work and articles, as well as interviews with experts in AI hardware and software, threat intelligence, engineering, and cybersecurity. Information on policies, recent news, and data on AI or data center trends was derived from news articles, think tank reports and briefs, and interviews with governance, cybersecurity policy, and foreign and domestic policy experts. In total, over 100 sources and five expert interviews contributed to this report's findings.

The Evolution of Data Centers

A Brief History

The prototypes of data centers have been around since the 1950s and 1960s and were called mainframes. Mainframes were single, isolated supercomputers that processed data and executed complex calculations. Mainframes were not connected to a network until the 1990s, when multiple microprocessor computers—later servers—started replacing mainframes.¹³ The modern data center is a physical facility with servers, networking equipment, storage devices, redundant power, and cooling infrastructure to store or process a large amount of data and to compute calculations. Data centers are critical information technology (IT) infrastructure because they store, distribute, and interpret data that is foundational to organizations’ day-to-day operations.¹⁴ As data centers get increasingly complex due to evolving technologies, operators deploy smart control systems and management software to optimize performance and energy efficiency.¹⁵ Figure 1 illustrates the increasing size of data centers to help envision the evolving complexity.

Figure 1 | Increasing Size of Data Centers



Source: “Data Centers: A Timeline of Growth and Expansion,” Datacate, accessed July 13, 2025, <https://www.datacate.net/data-centers-a-timeline-of-growth-and-expansion/>. Graphics by Alex Briñas/New America.

Data Centers

The Telecommunications Industry Association (TIA) has a system that classifies data centers into four tiers based on data center design—including the center’s architecture and topology, environmental design, power and

cooling systems, cabling systems, redundancy, and safety and physical security—which affects resiliency. Tier 4 data centers are the most resilient, and Tier 1 centers are the least resilient as represented in Figure 2 below.¹⁶

Figure 2 | TIA's Data Center Tiers

	Tier 1	Tier 2	Tier 3	Tier 4
Can operate through outages and power spikes due to Uninterruptible Power Supply (UPS)	✓	✓	✓	✓
IT systems physically inside data center	✓	✓	✓	✓
24/7 cooling system and backup generator	✓	✓	✓	✓
Can operate through human error-caused disruptions	✓	✓	✓	✓
Additional cooling (engine generators, cooling units, chillers, pumps, etc.)		✓	✓	✓
Can operate through maintenance and component removals		✓	✓	✓
Power and cooling redundancy support			✓	✓
Can operate through expected and unexpected disruptions due to physically isolated systems				✓
Fully redundant systems and fault-tolerant				✓
 Annual downtime	29 hours	22 hours	1.6 hours	26 minutes

Source: Graphic by Alex Briñas/New America.

In addition to the tier system, latency also factors into data center performance. Latency refers to the time it takes for the data center to receive and process a user request. Low latency signifies high performance with less delay and time spent before the data center responds to the request.¹⁷ While Tier 4 data centers with the lowest latency possible may sound like the goal for all data centers, not all centers require this level of capability because they vary in their specific purposes and critical requirements. To highlight this point, Table 1 provides a high-level summary of different types of data centers and a sample set of critical requirements.¹⁸

Table 1 | Examples of Types of Data Centers and Relevant Requirements

Type of Data Center	Description	Example of Requirements
Enterprise	Privately owned data centers for mission-critical applications at banks, health care systems, and other organizations' internal operations	<ul style="list-style-type: none"> • High levels of security • Highly fault-tolerant with redundant systems
Cloud	Third-party-owned and managed data centers that provide scalable, cloud-based services	<ul style="list-style-type: none"> • Low latency • Minimal downtime
Hyperscalers	Larger-scale cloud data centers that support large-scale applications and services, such as Amazon Web Services (AWS), Microsoft Azure, Google Cloud Platform, IBM Cloud, and Oracle	<ul style="list-style-type: none"> • Low latency • High compute power and resources • High redundancy
Edge	Smaller data centers placed near users in small cities and typically used for real-time data processing, such as for autonomous vehicles or real-time analytics	<ul style="list-style-type: none"> • Minimal latency
Telecom	Data centers that support telecommunications network functions and enable communication services	<ul style="list-style-type: none"> • High redundancy
Crypto mining	Data centers that specialize in crypto mining processes	<ul style="list-style-type: none"> • Low latency • Cheap and continuous electricity
IT system	Basic data centers that support everyday office needs	<ul style="list-style-type: none"> • Basic capacities • Could be a Tier 1 data center

NEW AMERICA

Given limited resources and the tradeoff between productivity and security, data centers prioritize different requirements when planning, constructing, and operating depending on their main purpose. Supporting AI is an emerging purpose for data centers. While there are still many unknowns about the related requirements, experts believe that AI data centers require high efficiency, significant power and electricity, high compute power, and low latency.¹⁹

AI Data Centers

An AI data center is commonly defined as an emerging type of data center built to support the high computational requirements of AI workloads. They can be divided into two subcategories: training AI data centers and inference AI data centers.²⁰ Training AI data centers process large amounts of data to train AI models and conduct machine learning training, while inference AI data centers deliver AI-driven insights and support the deployment of AI models into applications for end users.²¹

Like traditional data centers, AI data centers have **five major components**: (1) compute resource, (2) data storage system, (3) network infrastructure, (4) power or energy capability, and (5) cooling system—but in the AI context, all five components need to demonstrate high performance and productivity.²² (Table A1 in the [appendix](#) summarizes the differences between AI and traditional data centers across these five major components.) Of these five components, this report focuses on the differences in compute resources due to the different impacts they have on AI data center cybersecurity.

Traditional data centers typically utilize central processing units (CPUs) that use logic circuitry to process data and execute commands. CPUs can have multiple cores that can support multiple software types, but they cannot process operations simultaneously. CPUs also do not have enough memory to support AI data centers' workloads.²³ Therefore, AI data centers use more powerful compute resources:

- **Graphics processing units (GPUs)** allow for parallel operations. They are often used for AI training data centers and sometimes in other types of large-scale data centers.
- **Field-programmable gate arrays (FPGAs)** and **application-specific integrated circuits (ASICs)** can be customized to efficiently process AI workloads.²⁴
- **Google's tensor processing units (TPUs)** are a type of ASIC that optimizes performance for specific frameworks and effectively supports deep learning.

A key tradeoff when opting for the more powerful compute resource, such as GPUs, is that they may require much more power and create more heat than

CPUs.²⁵ Also, while CPUs and the more powerful compute hardware options face similar cyber threats, such as side-channel attacks, GPUs are vulnerable to additional cyber threats, which will be further detailed in the **AI Data Center Threats** section.

Beyond the five major components, supplementary considerations exist, such as facility location. As of April 2024, about 50 percent of the over 2,500 data centers in the United States are in Northern Virginia, Northern California, and Dallas, Texas, to prioritize low latency.²⁶ Yet new AI data centers, especially training ones, are now being constructed in more remote areas—for example, in Indiana, Iowa, and Wyoming—to be closer to power plants and farther from cities due to concerns of draining power grids.²⁷ AI inference data centers can be closer to cities and users: Cerebras AI is planning to build them in Santa Clara, California; Atlanta, Georgia; and Montreal, Canada.²⁸ The location of AI data centers—which often depends on latency and power—can increase risks, especially if AI data centers are built in countries with cheap energy and land, as in Southeast Asia.²⁹ This risk will be further detailed in the **Changing World Threats** section.

Given that AI data centers require the same five components as traditional ones, the two types of centers face similar cybersecurity threats. However, AI data centers face unique or increased threats due to different hardware requirements and location. Additional threats from evolving technologies and changing geopolitics also heighten security requirements.

Cyber Threats

The different pieces of a data center—the purposes or types, the five key components, and additional considerations—influence a data center’s vulnerabilities and risks, and all need to be secured for a cyber-secure AI data center. This includes the requirements of AI workloads: the models, weights, and training data.

Additionally, while not a foundational requirement, increasingly complex and large data centers utilize data center infrastructure management (DCIM) software to efficiently maintain and manage the facilities. DCIMs will be useful for AI data centers in the following capacities:³⁰

- **Intelligent capacity search** enables quickly finding which device or rack has the capacity for additional AI deployment.
- **Predictive analysis** predicts the impact of AI workloads on rack capacity and power usage.
- **Automatic server power budgeting** calculates the required power for servers.
- **Dynamic single-line power diagrams** depict power capacity and load to support redundancy planning and protect against break trips.

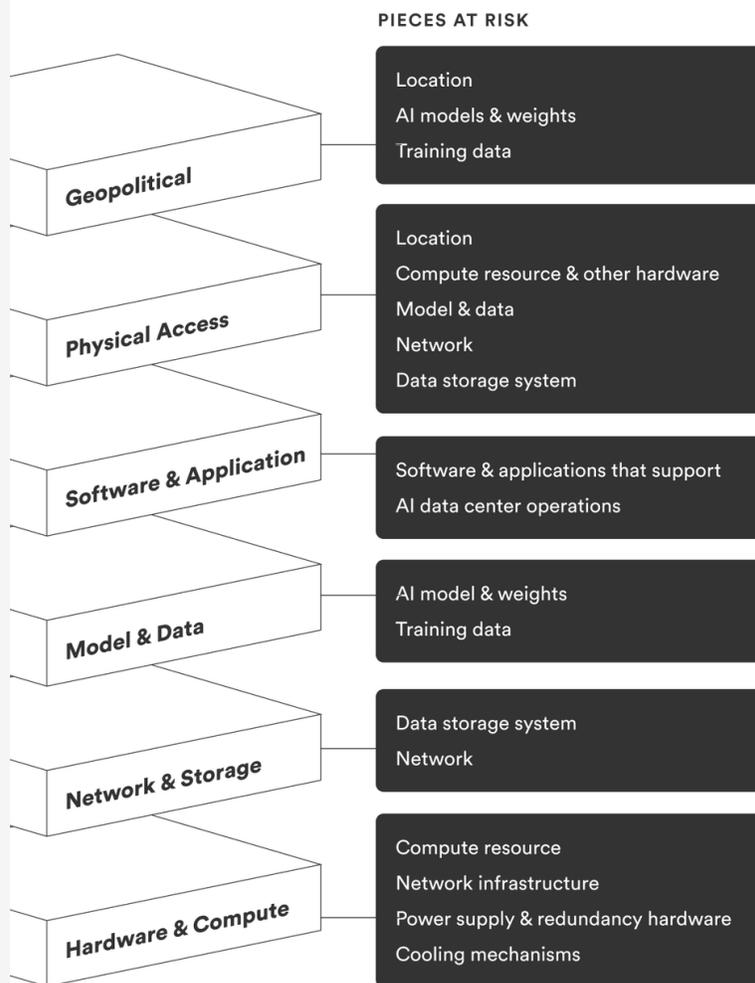
There is also newer software designed specifically for certain hardware mostly found in AI data centers, such as Trend Vision One—Sovereign and Private Cloud (SPC) for GPUs and TensorFlow for TPUs, further emphasizing the importance of cyber-secure software and applications deployed in AI data centers.³¹

→ A SECURITY FRAMEWORK FOR AI DATA CENTERS

In considering all the pieces inside an AI data center, this report proposes the following framework with six layers of security: (1) hardware & compute, (2) network & storage, (3) model & data, (4) software & application, (5) physical access, and (6) geopolitical.

Some elements may require cybersecurity considerations at multiple layers depending on the threats they face. For example, training data in AI data centers requires security at the model & data layer from cyberattacks, at the physical access layer from infiltrators who physically enter the facilities to steal the data, and at the geopolitical layer from state-sponsored actors who are targeting AI training data for their own AI sovereignty. Figure 3A below maps the data center pieces at risk at each of the security layers.

Figure 3A | Six Layers of AI Data Center Security



This section of the report will dive into the different threats faced by both traditional and AI data centers and map them to the six-layer framework.

Traditional Data Center Threats

Data centers are high-value targets because they can store important data, provide critical services, and run networks, applications, security, and virtual machines. As previously mentioned in Table 1, data centers support mission-critical applications at banks and health care systems, allow cloud-based services to be scalable, process data for autonomous vehicles and the Internet of Things, and enable communication services. It is not an understatement to say that data centers are the backbone of all online services. Thus, a disruption or outage at data centers—whether caused intentionally by threat actors or unintentionally by human error or process failures—can be costly for data center operators, risking financial loss, reputation damage, regulatory penalties, and loss of customers.³² For example, in 2021, when data center outages cost companies an average of \$100,000 per incident, Amazon Web Services experienced a five-hour data center outage due to a failure in network devices in the US-EAST-1 Region, costing the company \$34 million in revenue.³³

Both sophisticated nation-state actors and financially motivated cyber criminals target data centers with disruption attacks, ransomware incidents, and data breaches.³⁴ Four common cyberattacks on data centers are listed in Table 2 along with examples of incidents.³⁵

Distributed denial of service (DDoS) attacks rely on network connectivity to be able to target and overload a system or service, making defense at the network & storage layer critical.³⁶ A ransomware attack targets and encrypts data, highlighting the importance of cybersecurity at the model & data layer.³⁷ A supply chain attack can make both hardware and software inside a data center vulnerable, proving hardware and software layers to be crucial.³⁸ Social engineering attacks depend on threat actors gaining a form of access into the data center whether through calling and deceiving the target's customer support or manipulating employees to install remote access Trojans (RAT)—malware that gives hackers remote access into and control of a targeted system—or through backdoors into data center components.³⁹ Social engineering attacks necessitate model & data as well as physical access layers of security. Furthermore, the 2017 DDoS campaign against Google and the 2018 Supermicro supply chain attack referenced in Table 2 involved actors in China, requiring consideration of the geopolitical layer that will be further detailed in the **Changing World Threats** section.

Table 2 | Four Common Cyber Attacks on Data Centers

Type of Attack	Description	Example Incident(s)
Distributed denial of service (DDoS)	Hinders digital services by sending too many requests to a web server and overwhelming internet bandwidth, CPU, and random-access memory (RAM) capacity	<ul style="list-style-type: none"> • 2017: six-month 2.54 Tbps DDoS campaign on Google’s servers that was attributed to Chinese hackers
Ransomware	Inhibits victims from accessing files, systems, and networks by encrypting the victims’ data and demanding a ransom in return for access restoration	<ul style="list-style-type: none"> • 2024: LockBit 3.0 ransomware attack on Indonesia’s Temporary National Data Center, in which the threat actors demanded an \$8 million ransom and the nation’s immigration services and education enrollment platforms were hindered
Supply chain	Compromises a data center by targeting a third-party vendor that provides hardware or services to the data center and therefore has access to the center’s systems, networks, and data	<ul style="list-style-type: none"> • 2018: Chinese supply chain attack targeting U.S. manufacturer Supermicro compromised data centers, potentially even those of the U.S. Department of Defense, by implanting networking monitoring and control chips on motherboards for Supermicro • 2023: application attack—gaining unauthorized access by exploiting code vulnerabilities (CVE-2023-22527) of deployed applications—on an Atlassian confluence data center and server that allowed threat actors to illicitly and covertly mine cryptocurrency
Social engineering	Tricking a victim into sharing sensitive information that can allow unauthorized access or assist with fraud	<ul style="list-style-type: none"> • 2021: Shanghai-based GDS Holdings Ltd and Singapore-based ST Telemedia Global Data Centers breach using passwords received from deceived customer support. Breach impacted over 2,000 organizations, with China’s main foreign currency and debt-trading site hacked as a result of the data center breach

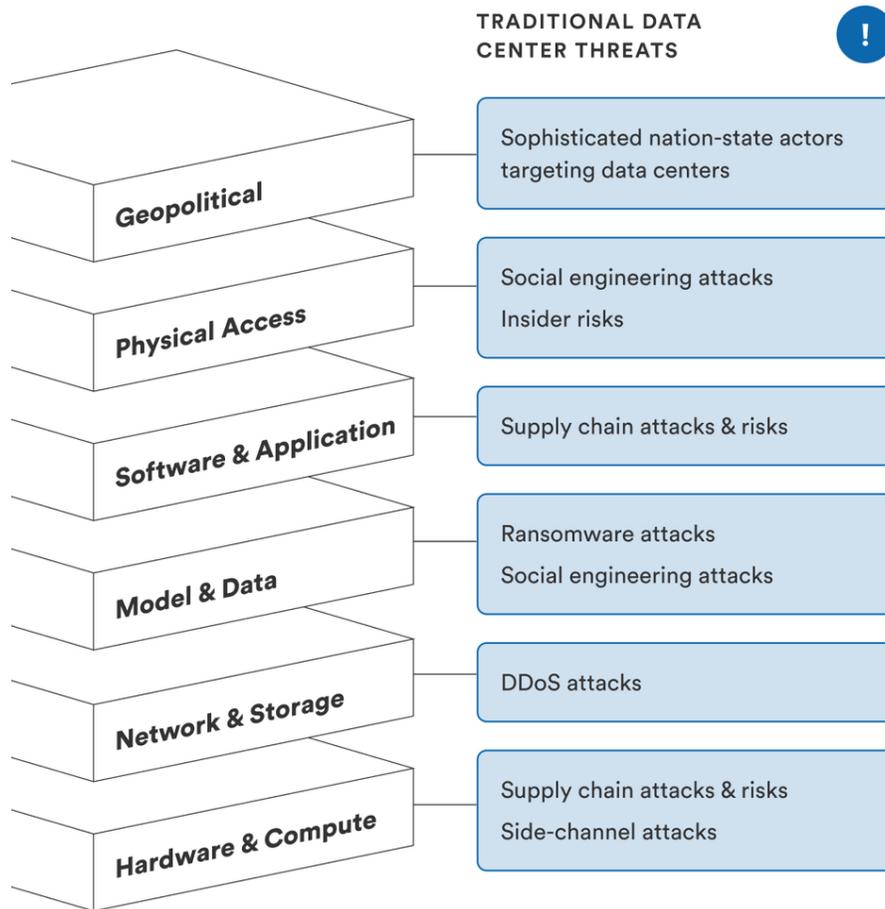
NEW AMERICA

Threat actors can also combine these types of attacks in one operation. For example, in 2024, the German data center power supply company Bender experienced a ransomware attack in which hackers infiltrated its operating systems and gained unauthorized access to account, financial, and banking data.⁴⁰ This incident, which affected the power and operations of a data center, was a combination of ransomware and supply chain attacks.

Beyond the four common types of cyberattack, there is another type of attack that data centers are particularly vulnerable to: side-channel attacks, a cybersecurity attack that collects information from a system's process and execution or attempts to influence a system's program.⁴¹ In data centers, side-channel attacks include measuring or surveilling fan power or the sounds a CPU makes, or using sensors that measure the electromagnetic field. These attacks can show threat actors CPU-level activity, the data architecture, and data usage, requiring security at the hardware & compute layer. In July 2025, global semiconductor company Advanced Micro Devices reported that four new processor vulnerabilities could allow hackers to conduct timing-based side-channel attacks, and cybersecurity company CrowdStrike labeled these weaknesses as critical threats.⁴²

Finally, there are two specific types of risks also critical to data center security. The first is insider risk, defined by the Cybersecurity and Infrastructure Security Agency as “the potential for an insider to use their authorized access or understanding of an organization to harm that organization,” whether intentional or unintentional.⁴³ Insider risk can be mitigated by limiting access to trusted people, and it requires cyber defense at the physical access layer. The second is supply chain risk. Beyond supply chain attacks described earlier in this section, there is an added risk because a majority of the components in a data center are developed by third-party suppliers. At the hardware layer, data centers rely on a global network of hardware suppliers, which creates risks of counterfeit and flawed hardware even without a threat actor's interception or attack.⁴⁴ At the software layer, third-party tools and applications hosted on data centers can have vulnerable code, making the selection of third-party suppliers as well as secure design of applications and code critical.⁴⁵ Figure 3B below maps traditional data center cyberattacks and risks to the proposed six layers of security.

Figure 3B | Traditional Data Center Threats



Source: Graphic by Alex Briñas/New America.

AI Data Center Threats

In addition to all the threats that traditional data centers face, AI data centers face unique threats. Specifically: (1) The **hardware & compute layer needs to consider ASIC and AI-specific hardware-level attacks**, such as memory-level and TPU-specific attacks, and (2) the **data layer needs to expand to include model security** since AI model weights and training data are in AI data centers.

At the hardware & compute layer, defending against supply chain attacks and vulnerabilities, as well as side-channel attacks, are still critical. AI data centers need trusted and secured GPUs and ASICs, which, like CPUs, rely on a global supply chain.⁴⁶

Once the hardware is inside the AI data center, the sounds from GPUs can reveal information to threat actors about those GPUs, model architecture, and architecture weights.⁴⁷ A common GPU side-channel attack is keystroke inference, through which hackers monitor GPU-based rendering workloads to get access to keystrokes and user inputs.⁴⁸ Tim Fist, director of emerging technology at the Institute for Progress, highlighted that “GPUs are more vulnerable than CPUs to memory-level attacks.”⁴⁹ GPUs do not always have sufficient memory isolation, which means that memory can be carried from one process to another and allow threat actors to get access to model weights and training data.⁵⁰ There is also GPU-specific malware that can execute malicious code on a GPU’s memory and could bypass traditional CPU security tools. Notably, in January 2025, Nvidia announced that seven new vulnerabilities were found in its GPUs, with three of them being of high severity.⁵¹

GPU vulnerabilities can be found in traditional data centers because large-scale centers deploy GPUs. On the other hand, TPUs—which are also subject to the side-channel and memory-level attacks mentioned above—are uniquely designed for AI and were first deployed internally by Google in 2015.⁵² In early 2025, a TPU-specific side-channel attack, TPUXtract, was discovered. TPUXtract exploits unintentional TPU data leaks to enable a threat actor to infer an AI model’s parameters, essentially allowing AI model exfiltration and intellectual property (IP) theft.⁵³

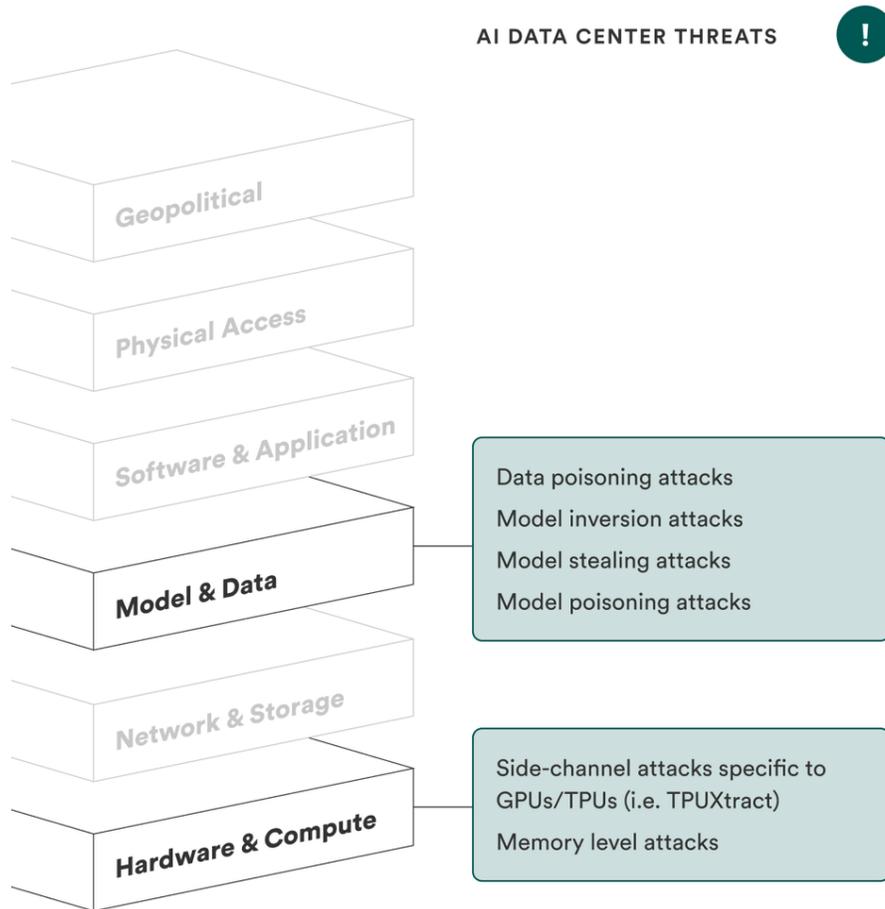
Given that the hardware & compute layer threats enable hackers to extract information about the AI model and weights deployed in the AI data center, the threats at the data layer can extend to the model layer. When data is exfiltrated from a traditional data center, the impacts can include the loss of sensitive data, regulatory issues, and reputational damage.⁵⁴ When AI model information and IP are accessed and stolen, the impacts expand to include risking the integrity and confidentiality of the AI model. Threat actors can manipulate exfiltrated models to create vulnerabilities and bias decision-making.⁵⁵ Additionally, as it becomes increasingly evident that AI will influence geopolitics, the global economy, national security, and human lives, the importance of protecting models becomes critical.⁵⁶ This point will be further described in the [Changing World Threats](#) section.

Model threats include AI training and model weight exfiltration through traditional ransomware and social engineering attacks. Furthermore, AI models face specific model-level threats:⁵⁷

- **Data poisoning attacks:** accidentally or purposefully including incorrect data in the AI training dataset, leading to erroneously trained AI models
- **Model inversion attacks:** recovering training data from AI models by querying models, examining the outputs, and extracting information
- **Model stealing attacks:** querying AI models and using the outputs to train a replacement model trained like the original model
- **Model poisoning attacks:** modifying model parameters or architecture to create a backdoor or change the model's behavior⁵⁸

At a high level, AI hardware-specific cyber threats and model threats expand cyber threats to AI data centers. Figure 3C below maps cyberattacks specific to AI data centers to the proposed six layers of security.

Figure 3C | AI Data Center Threats



Source: Graphic by Alex Briñas/New America.

Changing World Threats

As data centers and their components change to meet the requirements of AI data centers, the world continues to transform as well, introducing new or enhanced threats for emerging AI data centers.

The first type of change is in volatile geopolitics. While traditional data centers also face geopolitical threats and are targeted by nation-state actors, these threats are heightened for AI data centers due to AI’s significance for national security and economic competitiveness. With the advent of the AI race, the concept of **sovereign AI**—the ability of a nation to develop, use, and govern its own AI models and related infrastructure—is on the rise, and many believe

that sovereign AI is crucial for national security and economic competitiveness.⁵⁹

In April 2025, the U.S. Department of Defense highlighted that the Joint Staff is implementing AI to improve military operations, to improve a commander's decision-making and responsiveness, and to streamline processes.⁶⁰ The potential for AI technologies to help counter threats across all sectors, including critical energy infrastructure, will only embed AI deeper into critical infrastructure.⁶¹ In terms of economic competitiveness, technological innovation results in high-level economic improvement, and with AI specifically, automation of routine tasks and AI-supported creative and technical work are predicted to further innovation, allow for strategic tasks, and improve a nation's economy.⁶²

These uses for AI indicate that if actors stole AI models, weights, or training data, sensitive data related to the military, economy, critical infrastructure, and more would be leaked. Additionally, stolen AI models can help hackers create deepfakes or convincing phishing emails for enhanced psychological warfare and disinformation campaigns, election interference, cybercrime and financial fraud, and critical infrastructure attacks.⁶³

As AI becomes a key component of national security and economic competitiveness, AI data centers become key assets to protect from foreign adversaries, especially from sophisticated cyber threat actors sponsored by Russia, China, Iran, and North Korea. Nation-state threat actors are well resourced and sophisticated, posing a serious threat to even the biggest U.S. companies and critical industries, as CrowdStrike's chief security officer Shawn Henry stated in 2023.⁶⁴ For example, the Chinese state-sponsored hacking group Salt Typhoon successfully gained access to an unprecedented amount of information from the largest U.S. telecommunications companies and accessed information on high-value targets like Donald Trump and JD Vance.⁶⁵

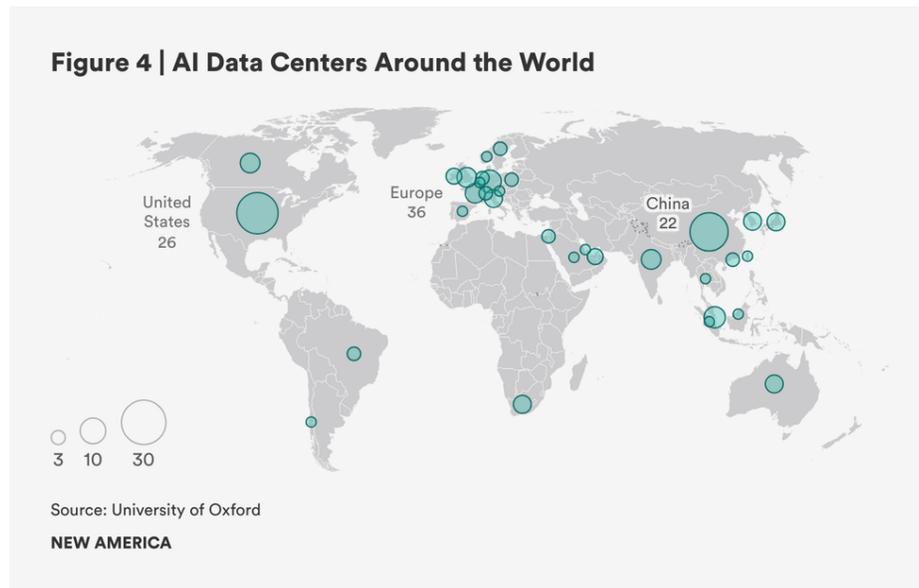
Unfortunately, like most businesses and companies, AI data centers' commercial operators are typically not equipped to defend against such sophisticated operations. In 2024, OpenAI emphasized the importance of increasing AI data center and infrastructure security.⁶⁶

Even before an AI data center is operational, there is potential for sabotage attempts, given that Chinese companies exclusively manufacture many AI data center components, such as most transformer substations critical for power systems. This means that Chinese companies could install backdoors into the hardware components. In addition, most AI-specific GPUs are made in Taiwan, which is threatened by China as part of a larger national sovereignty debate. Historically, Taiwan Semiconductor Manufacturing Company (TSMC) technologies and products have been unlawfully transferred to mainland

China. There are also concerns and a strong suspicion that the Chinese Communist Party (CCP) has infiltrated TSMC and U.S. labs with spies.⁶⁷

Once an AI data center is operational, geopolitical threats continue as AI data centers are vulnerable to state-sponsored disruption and exfiltration attacks due to an insufficient level of cybersecurity.⁶⁸ For example, Russia has sophisticated cyber capabilities to infiltrate AI data centers, steal models, and potentially run a model in its own infrastructure.⁶⁹

Despite the potential state-sponsored targeting of AI data centers, companies and nations are risking increased threats by building AI data centers abroad and joint centers shared with other countries. Figure 4 below shows the number of AI data centers in different countries, and also reveals that Asia has the most.⁷⁰ In January 2025, the United States and the United Arab Emirates (UAE) announced a partnership to build a data center in Abu Dhabi, which will be the largest AI data center outside the United States.⁷¹ Concerningly, the Persian Gulf is a part of China's Digital Silk Road 2.0, and the UAE has adopted Chinese 5G technology and city-wide surveillance programs, potentially giving the CCP access to the AI data center. Companies also choose to build AI data centers abroad in locations that have cheap energy and land, such as Malaysia.⁷²



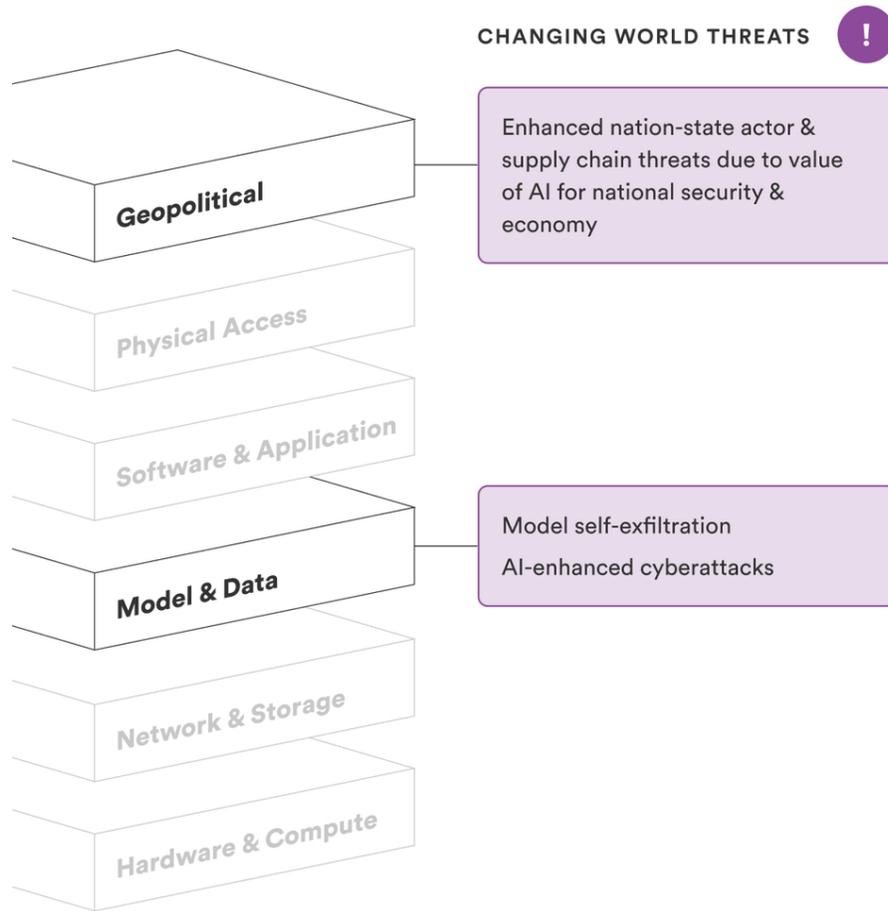
The second type of change is evolving technologies. As previously mentioned, the capabilities of AI data centers have conversely enabled an increasingly robust system, which wielded by U.S. competitors, increases risks to critical data center infrastructures. Threat actors are using AI tools to enhance phishing campaigns, research target networks, conduct post-compromise

activities, and support coding tasks.⁷³ They are also using AI-generated deepfakes to bypass multifactor authentication.⁷⁴

AI models themselves can also pose a threat to data centers through AI self-exfiltration. Traditional data centers face threats from backdoors installed into data center components, allowing a threat actor to exfiltrate or access sensitive information. AI data centers face an additional threat of backdoors installed into AI models through model poisoning, which could enable threat actors to steal the model weights and information. In addition, model self-exfiltration is a novel threat that refers to AI models deceiving the user and cybersecurity measures in place in order to self-leak its model weight and sensitive data.⁷⁵ Typical methods to protect model weights—such as air gaps—may not stop model self-exfiltration because the model may take over the system.⁷⁶ This new threat—a novel form of data exfiltration and leakage—further necessitates enhanced model security.

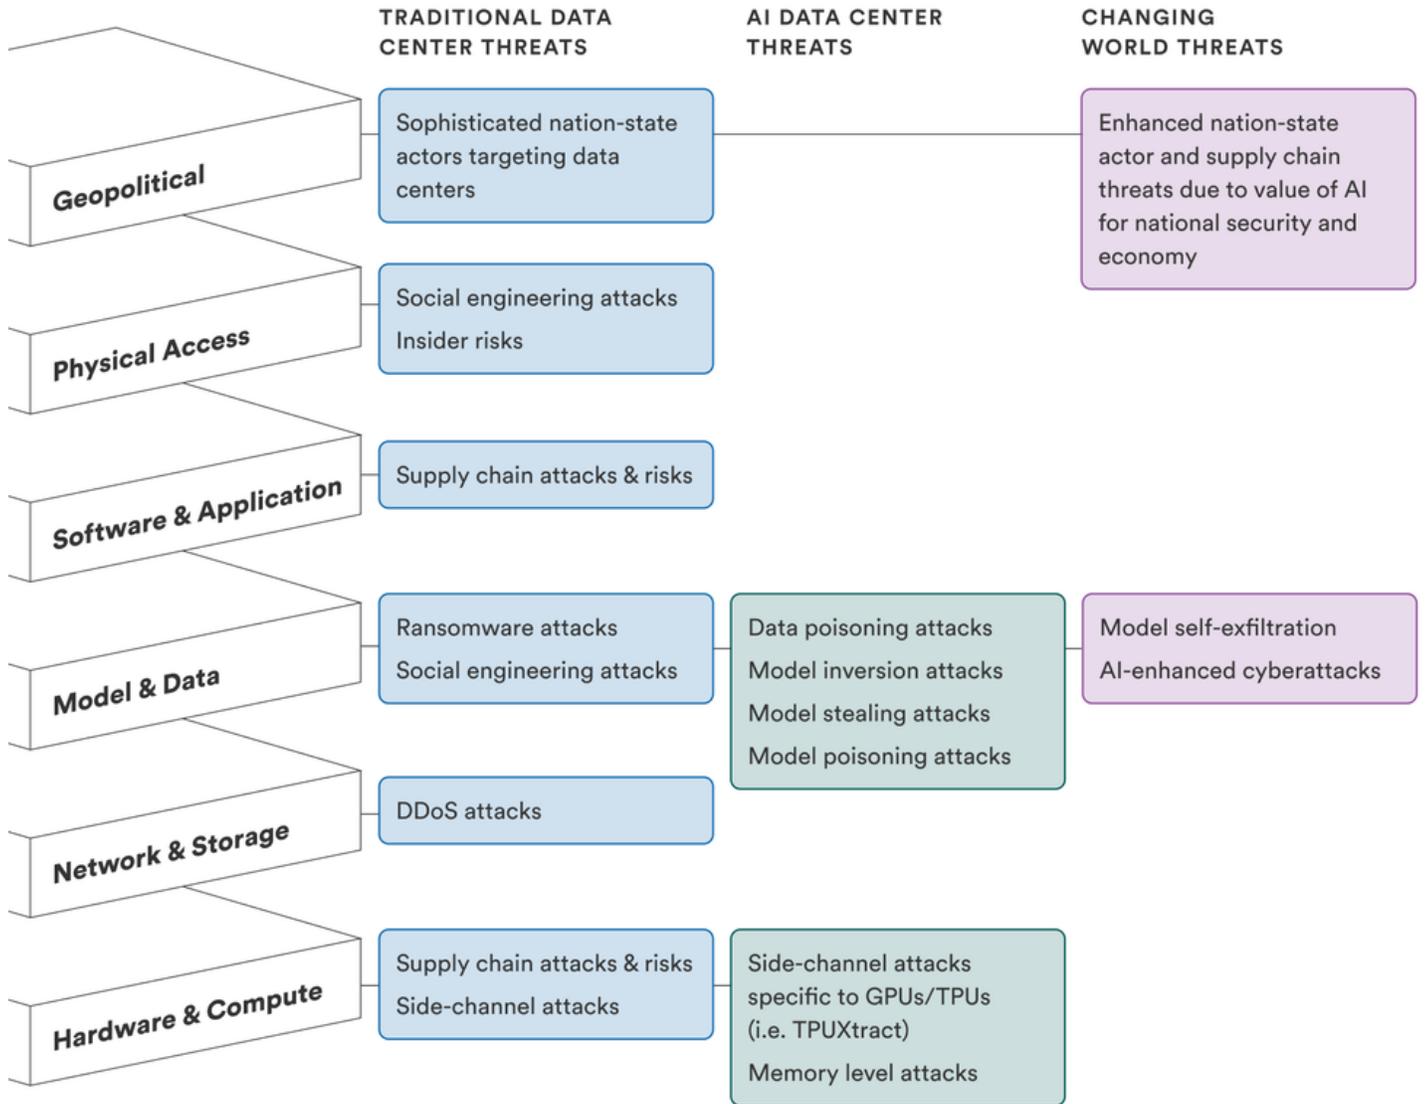
Besides AI, there are other technological evolutions in the works that will enhance threats to AI data centers, such as quantum computing, which will soon require next-generation cryptography.⁷⁷ Figure 3D below maps changing world threats to the proposed six layers of security, and Figure 3E compiles Figures 3B–3D to map the different types of threats across the layers.

Figure 3D | AI Data Center Threats



Source: Graphic by Alex Briñas/New America.

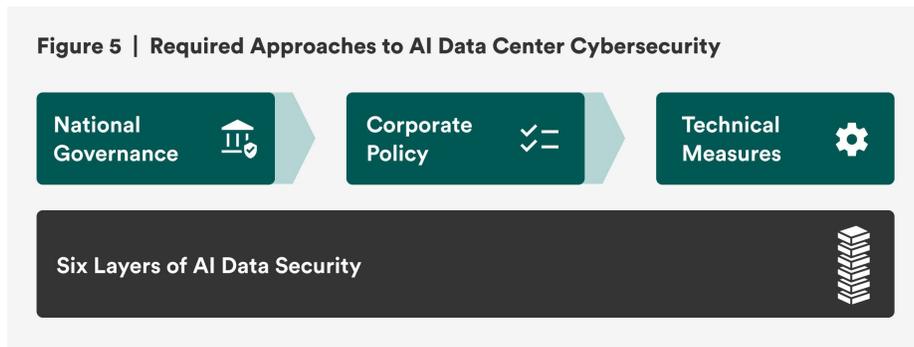
Figure 3E | Overview of Threats



Source: Graphic by Alex Briñas/New America.

A Framework for Cyber-Secure AI Data Centers

Given the threats and high-value nature of the target, the cybersecurity of AI data centers needs to be top tier. As Tim Fist commented, “AI data centers used to train and run the most powerful models will likely need to be secured with nation-state-level adversaries in mind—this will likely require taking some of the measures typically only used on government data centers used to store and process highly classified information, as well as many additional measures specific to AI performance and security requirements.”⁷⁸ In line with this comment, this report suggests that each of the six layers of security require three approaches: technical, corporate policy, and national governance (see Figure 5). Bridge the gaps between the technical, corporate policy, and national governance approaches with a framework that maps the threats to AI data centers across the six layers of security.



Source: Graphic by Alex Briñas/New America.

Technical Measures

Recommendation 1: Implement existing research and standards for technical requirements in AI data centers.

Informing the security approach from a technical perspective, RAND published a report in May 2024 that created security levels from one to five for securing model weights.⁷⁹ Security Level 1 (SL1) indicates an AI system that can defend against amateur attacks, and Security Level 5 (SL5) can protect against the most sophisticated attacks—even ones by nation-state actors. The report includes a benchmark for each level, detailing technical security measures necessary for that level.⁸⁰ The Institute for Progress built on RAND’s work and outlined an overview of technical measures required for AI data centers to reach SL4 across

supply chain, network & storage, hardware, and physical access security.⁸¹ The technical needs for a cyber-secure AI data center are thus generally well known. However, corporate policy and national governance measures need to be in place for a top-tier consolidated cybersecurity approach that incentivizes AI data center companies and operators to implement the required technical measures.

Corporate Policy Measures

Recommendation 2: Corporate policies need to require technical measures across the six layers of security.

For AI data center companies and operators to have the appropriate technical measures in place, they would need corporate governance measures that require the technical mitigations in a structured process. For example, corporate policies that require all AI data centers to have a Faraday cage or shield chamber deployed with the hardware are necessary in order to mitigate side-channel attacks on data center hardware and defend against tracking and monitoring of electromagnetic emanations.⁸² At the model & data layer, a corporate policy needs to require continuous AI audits and monitoring in order to secure AI models, identify backdoors and vulnerabilities in them, and defend against model and data exfiltration.

National Governance Measures

Recommendation 3: National governance measures should focus on incentivizing operators to meet the technical and corporate policies needed for a cyber-secure AI data center.

National governance approaches need to provide a framework that encourages responsible AI data center construction and operations. Even though companies have natural incentives—such as reputational risks from data leakage or financial risks from AI data center outages and disruptions—to make AI data centers cyber secure, defending these assets from the most sophisticated threat actors is not easy and requires significant resources and investments. National policy approaches are currently immature: While there are over 200 national or supranational AI laws and regulations, very few actually govern AI usage, deployment, and infrastructure with binding

legislation.⁸³ Some existing standards that can help support AI data center security include:⁸⁴

- The National Institute of Standards and Technology (NIST)'s Secure Software Development Framework (SSDF) includes best practices for decreasing software-level vulnerabilities and has an AI addendum that includes best practices for AI models.⁸⁵
- NIST's SP 800-171 provides standards for protecting unclassified information.⁸⁶
- NIST's SP 800-53 includes standards for security and privacy control.⁸⁷
- NIST's FIPS 140-3 outlines the design and operation of computer hardware that processes and protects sensitive data.⁸⁸
- The Federal Risk and Authorization Management Program (FedRAMP), based on NIST SP 800-3, looks at security assessments, authorization, and monitoring to determine whether a cloud service provider is in compliance with the program's standards.⁸⁹
- The U.S. Department of Defense has a Cybersecurity Maturity Model Certification (CMMC) program that requires defense contractors to have sufficient security measures to protect unclassified and sensitive information.⁹⁰
- The Cybersecurity and Infrastructure Security Agency (CISA) has a Zero Trust Maturity model that helps define best practices for controlling access to sensitive data.⁹¹
- CISA's Software Bill of Materials (SBOM) offers an ingredients list for software to identify software and supply chain vulnerabilities.⁹²

Policies and objectives focused on AI data centers have also emerged but have not consistently required nor incentivized AI data center security. President Biden's 2025 Executive Order (EO) 14141: Advancing United States Leadership in Artificial Intelligence Infrastructure required AI data center operators to submit security proposals if requesting to build on federal land.⁹³

Unfortunately, President Trump revoked EO 14141 in July 2025. Furthermore, the Trump administration's EO 14179: Removing Barriers to American Leadership in Artificial Intelligence and EO 14154: Unleashing American Energy led the Department of Energy to designate 16 potential federal sites for rapid AI data center construction that may not take into account security requirements.⁹⁴ AI data center construction on federal land has yet to materialize.

More recently, in July 2025, President Trump revealed his AI Action Plan, which further mandates that federal land be available for AI data centers and supporting infrastructure construction without security requisites.⁹⁵ Additionally, the action plan's central theme of "Build, Baby, Build!" encourages businesses to build AI tech stacks and data centers abroad and only mentions high-security technical standards for AI data centers utilized by the military and intelligence community.⁹⁶

National regulations with significant fines and risk-based frameworks that require stronger security measures in AI data centers are currently lacking but are crucial for incentivizing AI data center operators and companies to meet high-security technical requirements.⁹⁷ Benefits such as tax breaks or access to federal land tied to strong security requirements can also encourage operators.

Once national governance measures exist, they can further incentivize AI data center businesses and operators by highlighting the return on investment when complying with regulations: When there is a distinction between noncompliant and compliant AI data centers, investors, customers, and potential employees will flock to the more cyber-secure centers.⁹⁸

For example, at the software & application layer, supply chain attacks and vulnerabilities pose risks to AI data centers. A technical approach to mitigate the risk would be to implement secure coding or test software with penetration testing and red teaming.⁹⁹ The necessary corporate policy would be to only allow tested software and applications into the company's AI data centers and to require audits of source code before deploying the software.¹⁰⁰ Finally, national governance measures that allow only AI data center operators who follow CISA's Secure by Design approach¹⁰¹ or implement software with SBOMs to be government contractors can incentivize AI data center companies to implement corporate policies that meet the technical requirements.¹⁰²

Conclusion

AI data centers may be similar to traditional ones in various ways, but they require a heightened threshold for security because they face expanded risks. With the global race to develop AI power and sovereign AI, securing AI and the infrastructure behind it—AI data centers—becomes critical for national security, the economy, defense, energy, and more. This report recommends that in order to develop cyber-secure AI data centers, there must be a framework that aligns the technical, corporate policy, and national governance approaches to cover the six layers of security: hardware & compute, network & storage, model & data, software & application, physical access, and geopolitical.

First, since the technical needs for a cyber-secure AI data center are known, operators must implement existing research and standards in an AI data center. Second, corporate policies need to require technical measures across the six layers of security. Finally, national governance measures should focus on incentivizing AI data center operators to meet the technical needs and implement corporate policies for a cyber-secure AI data center.

Appendix

Table A1 | Differences Between Traditional & AI Data Centers

	Traditional Data Center	AI Data Center
Compute Hardware	Serial Processing and Less Compute <ul style="list-style-type: none"> • CPUs 	Parallel Processing and High Compute <ul style="list-style-type: none"> • GPUs, FPGAs, ASICs, TPUs
Data Storage	Traditional for Structured Data <ul style="list-style-type: none"> • Hard disk drives (HDDs) • Solid-state drives (SSDs) • Network-attached storage (NAS) • Direct-attached storage (DAS) 	High-Performance for Unstructured Data <ul style="list-style-type: none"> • Non-volatile memory express (NVMe) Flash • Parallel file systems • Tiered storage
Network Infrastructure	Typically Moderate Bandwidth & Speed <ul style="list-style-type: none"> • Standard Ethernet • Average fiber optics 	High Bandwidth & Speed <ul style="list-style-type: none"> • Software-defined networking (SDN) • Low-latency fabrics • High bandwidth Ethernet or Infinibands
Power & Energy	Increasing <ul style="list-style-type: none"> • Often 5kW-10kW per rack • Typical PDUs • 12-volt servers 	High & Further Increasing <ul style="list-style-type: none"> • Two to four times more power than traditional centers • 15kW-60kW per rack • Larger Power Distribution Units (PDUs) • 48-volt servers
Cooling System	Traditional <ul style="list-style-type: none"> • Cold air circulation 	Innovative and Efficient <ul style="list-style-type: none"> • Liquid cooling & immersion cooling • Direct-to-chip (DTC) technology

NEW AMERICA

Notes

- 1 “Data Centers,” Data Center Map, accessed July 10, 2025, <https://www.datacentermap.com/datacenters/>; Bhargs Srivathsan et al., “AI Power: Expanding Data Center Capacity to Meet Growing Demand,” McKinsey & Company, October 29, 2024, <https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/ai-power-expanding-data-center-capacity-to-meet-growing-demand>.
- 2 Ryan Pendell, “AI Use at Work Has Nearly Doubled in Two Years,” Gallup, June 16, 2025, <https://www.gallup.com/workplace/691643/work-nearly-doubled-two-years.aspx>; Adam Satariano and Paul Mozur, “The Global AI Divide,” *New York Times*, June 21, 2025, <https://www.nytimes.com/interactive/2025/06/23/technology/ai-computing-global-divide.html>.
- 3 Deepa Seetharaman and Tom Dotan, “Tech Leaders Pledge Up to \$500 Billion in AI Investment in U.S.,” *Wall Street Journal*, January 21, 2025, <https://www.wsj.com/tech/ai/tech-leaders-pledge-up-to-500-billion-in-ai-investment-in-u-s-da506cd4>; Meghan Bobrowsky, “Meta Spending to Soar on AI, Massive Data Center,” *Wall Street Journal*, January 25, 2025, <https://www.wsj.com/tech/ai/meta-spending-ai-facebook-data-centers-9452a88f>.
- 4 École Polytechnique, “MGX, Bpifrance, Mistral AI, and NVIDIA Launch Joint Venture to Build Europe’s Largest AI Campus in France,” May 19, 2025, <https://www.polytechnique.edu/en/press-room/press-releases/mgx-bpifrance-mistral-ai-and-nvidia-launch-joint-venture-build-europes-largest-ai-campus-france>; Amber Jackson, “New U.S. \$2.3 bn AI Data Centre by EDGNEX Hailed a ‘Milestone,’” *Data Centre Magazine*, June 17, 2025, <https://datacentremagazine.com/critical-environments/new-us-2-3bn-ai-data-centre-by-edgnex-ailed-a-milestone/>; K. Oanh Ha, Yang Yang, and Naomi Garyan Ng, “China’s Got Big Plans for AI—In the Desert,” *Bloomberg*, July 8, 2025, <https://www.bloomberg.com/news/articles/2025-07-08/china-builds-ai-dreams-with-giant-data-centers-in-the-desert>.
- 5 Tim Fist and Arnab Datta, *How to Build the Future of AI in the United States* (Institute for Progress, October 23, 2024), <https://ifp.org/future-of-ai-compute/>.
- 6 Beth Stackpole, “AI Has High Data Center Energy Costs—But There Are Solutions,” MIT Sloan, January 7, 2025, <https://mitsloan.mit.edu/ideas-made-to-matter/ai-has-high-data-center-energy-costs-there-are-solutions>.
- 7 “AI Is Set to Drive Surging Electricity Demand from Data Centres While Offering the Potential to Transform How the Energy Sector Works,” International Energy Agency, April 10, 2025, <https://www.iea.org/news/ai-is-set-to-drive-surging-electricity-demand-from-data-centres-while-offering-the-potential-to-transform-how-the-energy-sector-works>.
- 8 Joseph R. Biden, *Executive Order on Advancing United States Leadership in Artificial Intelligence Infrastructure*, 88 Fed. Reg. 10,001 (White House Archives, January 14, 2025), <https://bidenwhitehouse.archives.gov/briefing-room/presidential-actions/2025/01/14/executive-order-on-advancing-united-states-leadership-in-artificial-intelligence-infrastructure/>.

- 9 At the time of drafting this report, President Donald Trump had signed an executive order that included a review of NSM-22: Donald J. Trump, *Achieving Efficiency Through State and Local Preparedness*, March 19, 2025, <https://www.whitehouse.gov/presidential-actions/2025/03/achieving-efficiency-through-state-and-local-preparedness/>; Joseph R. Biden, *National Security Memorandum on Critical Infrastructure Security and Resilience*, National Security Memorandum/NSM-22, April 30, 2024, <https://bidenwhitehouse.archives.gov/briefing-room/presidential-actions/2024/04/30/national-security-memorandum-on-critical-infrastructure-security-and-resilience/>.
- 10 Lester Ross, “China Rolls Out Critical Information Infrastructure Security Protection Regulations,” WilmerHale, August 19, 2021, <https://www.wilmerhale.com/en/insights/client-alerts/20210819-china-rolls-out-critical-information-infrastructure-security-protection-regulations>.
- 11 “Risks of Data Exfiltration,” SentinelOne, accessed July 15, 2025, <https://www.sentinelone.com/cybersecurity-101/cybersecurity/data-exfiltration/#risks-of-data-exfiltration>.
- 12 Riscure Security Solutions Team, “Case Study: How TPuXtract Leveraged Keysight Tools for AI Model Extraction,” Keysight, March 19, 2025, <https://www.keysight.com/blogs/en/tech/nwvs/2025/03/19/case-study-how-tpuextract-leveraged-keysight-tools-for-ai-model-extraction>.
- 13 “A Brief History of Data Centers,” Digital Realty, accessed June 14, 2025, <https://www.digitalrealty.com/resources/articles/a-brief-history-of-data-centers>.
- 14 “What Is a Data Center?,” Zscaler, accessed July 14, 2025, <https://www.zscaler.com/zpedia/what-is-data-center>.
- 15 Emil Sayegh, “The Billion-Dollar AI Gamble: Data Centers as the New High-Stakes Game,” *Forbes*, September 30, 2024, <https://www.forbes.com/sites/emilsayegh/2024/09/30/the-billion-dollar-ai-gamble-data-centers-as-the-new-high-stakes-game/>.
- 16 “What Are Data Center Tiers?,” Digital Realty, accessed June 14, 2025, <https://www.digitalrealty.com/resources/articles/introduction-to-data-center-tiers>; “What Is a Data Center?,” Amazon Web Services, accessed June 14, 2025, <https://aws.amazon.com/what-is/data-center/>.
- 17 “What Is Low Latency?,” Cisco, accessed June 15, 2025, <https://www.cisco.com/c/en/us/solutions/data-center/data-center-networking/what-is-low-latency.html>.
- 18 Sayegh, “The Billion-Dollar AI Gamble: Data Centers as the New High-Stakes Game,” <https://www.forbes.com/sites/emilsayegh/2024/09/30/the-billion-dollar-ai-gamble-data-centers-as-the-new-high-stakes-game/>; Melissa Palmer, “Hyperscalers: The Complete Guide to What, Why and How,” *Solar Winds* (blog), January 24, 2023, <https://www.solarwinds.com/blog/hyperscalers-the-complete-guide>.
- 19 Srivathsan et al., “AI Power,” <https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/ai-power-expanding-data-center-capacity-to-meet-growing-demand>; Sayegh, “The Billion-Dollar AI Gamble,” <https://www.forbes.com/sites/emilsayegh/2024/09/30/the-billion-dollar-ai-gamble-data-centers-as-the-new-high-stakes-game/>.
- 20 “AI Data Center,” Sunbird, accessed June 15, 2025, <https://www.sunbirdcim.com/glossary/ai-data-center>.
- 21 Sayegh, “The Billion-Dollar AI Gamble,” <https://www.forbes.com/sites/emilsayegh/2024/09/30/the-billion-dollar-ai-gamble-data-centers-as-the-new-high-stakes-game/>.

- 22 “What Is an AI Data Centre, and How Does It Work?,” Macquarie Data Centres, July 15, 2024, <https://www.macquariedatacentres.com/blog/what-is-an-ai-data-centre-and-how-does-it-work/>.
- 23 Jacob Roundy, “How Do CPU, GPU, and DPU Differ from One Another?,” TechTarget, February 5, 2025, <https://www.techtarget.com/searchdatacenter/tip/How-do-CPU-GPU-and-DPU-differ-from-one-another>.
- 24 “How Are AI Data Centers Changing Infrastructure?,” TRG Datacenters, accessed June 15, 2025, <https://www.trgdatacenters.com/resource/how-are-ai-data-centers-changing-infrastructure/>.
- 25 Brian Venturo, “The Redesign of the Data Center Has Already Started. Here’s What It Looks Like,” CoreWeave, March 19, 2024, <https://www.coreweave.com/blog/the-redesign-of-the-data-center-has-already-started>.
- 26 Mary Zhang, “United States Data Centers: Top 10 Locations in the USA” Dgtl Infra, April 11, 2024, <https://dgtlinfra.com/united-states-data-centers/>.
- 27 Srivathsan et al., “AI Power,” <https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/ai-power-expanding-data-center-capacity-to-meet-growing-demand>.
- 28 David Chernicoff and Matt Vincent, “Cerebras Unveils Six Data Centers to Meet Accelerating Demand for AI Inference at Scale,” *Data Center Frontier*, March 18, 2025, <https://www.datacenterfrontier.com/hyperscale/article/55273769/cerebras-unveils-six-data-centers-to-meet-accelerating-demand-for-ai-inference-at-scale>.
- 29 Dylan Butts, “Malaysia Is Emerging as a Data Center Powerhouse amid Booming Demand from AI,” CNBC, June 16, 2024, <https://www.cnbc.com/2024/06/17/malaysia-emerges-as-asian-data-center-powerhouse-amid-booming-demand.html>.
- 30 “Data Center Infrastructure Management,” Sunbird, accessed June 16, 2025, <https://www.sunbirdcim.com/glossary/data-center-infrastructure-management-dcim>.
- 31 Agam Shah, “Trend Micro, Nvidia Partner to Secure AI Data Centers,” *DarkReading*, June 6, 2024, <https://www.darkreading.com/cloud-security/trend-micro-nvidia-partner-to-secure-ai-data-centers>.
- 32 “Data Center Threats and Vulnerabilities,” Check Point Software Technologies, accessed June 18, 2025, <https://www.checkpoint.com/cyber-hub/cyber-security/what-is-data-center/data-center-threats-and-vulnerabilities/>.
- 33 Rich Miller, “Problems With AWS Network Devices Caused Widespread Cloud Outage,” *Data Center Frontier*, December 8, 2021, <https://www.datacenterfrontier.com/cloud/article/11427750/problems-with-aws-network-devices-caused-widespread-cloud-outage>; Bill Kleyman, “The Data Center Ransomware Attack That Costs You Everything,” *Data Center Knowledge*, September 1, 2023, <https://www.datacenterknowledge.com/data-breaches/the-data-center-ransomware-attack-that-costs-you-everything>.
- 34 Beth Maundrill, “Cybersecurity Implications of Data Centres as Critical National Infrastructure,” Infosecurity Europe, October 28, 2024, <https://www.infosecurityeurope.com/en-gb/blog/regulation-and-policy/cybersecurity-of-data-centres-uk.html>.
- 35 “Datacenter Vulnerabilities: 7 Life-Changing Attacks You Must Know,” Enterprise Engineering Solutions Corporation, accessed July 15, 2025, <https://www.eescorporation.com/datacenter-vulnerabilities/>; Catalin Cimpanu, “Google Says It Mitigated a 2.54 Tbps DDoS Attack in 2017, Largest Known to Date,” *ZDNet*, October 16, 2020, <https://www.zdnet.com/article/google-says-it-mitigated-a-2-54-tbps-ddos-attack-in-2017-largest-known-to-date/>; Daryna Antoniuk, “Indonesia’s National Data Centre Encrypted With LockBit Ransomware Variant,”

The Record, June 24, 2024, <https://therecord.media/indonesia-national-data-centre-hacked>;
Curtis Franklin, “Report: In Huge Hack, Chinese Manufacturer Sneaks Backdoors Onto Motherboards,” *DarkReading*, October 5, 2018, <https://www.darkreading.com/cyberattacks-data-breaches/report-in-huge-hack-chinese-manufacturer-sneaks-backdoors-onto-motherboards>; “Application Attacks,” Contrast Security, accessed July 15, 2025, <https://www.contrastsecurity.com/glossary/application-attacks>; Abdelrahman Esmail, “Cryptojacking via CVE-2023-22527: Dissecting a Full-Scale Cryptomining Ecosystem,” Trend Micro, August 28, 2024, https://www.trendmicro.com/en_us/research/24/h/cve-2023-22527-cryptomining.html; “Cyber Attacks on Data Center Organizations,” Resecurity, February 20, 2023, <https://www.resecurity.com/ar/blog/article/cyber-attacks-on-data-center-organizations>.

36 Josh Fruhlinger and Lucian Constantin, “DDoS Attacks: Definition, Examples and Techniques,” CSO, May 17, 2024, <https://www.csoonline.com/article/571981/ddos-attacks-definition-examples-and-techniques.html>.

37 Kurt Baker, “Introduction to Ransomware,” CrowdStrike, March 4, 2025, <https://www.crowdstrike.com/en-us/cybersecurity-101/ransomware/>.

38 “Third-Party Data Breaches: What You Need to Know,” Mitrastech, January 7, 2025, <https://mitrastech.com/resource-hub/blog/third-party-data-breaches/>.

39 Computer Security Resource Center, “Social Engineering,” National Institute of Standards and Technology, accessed July 15, 2025, https://csrc.nist.gov/glossary/term/social_engineering; “What Is a Remote Access Trojan?,” Fortinet, accessed July 15, 2025, <https://www.fortinet.com/resources/cyberglossary/remote-access-trojan>; “Data Center Threats and Vulnerabilities,” Check Point

Software Technologies, accessed June 18, 2025, <https://www.checkpoint.com/cyber-hub/cyber-security/what-is-data-center/data-center-threats-and-vulnerabilities/>.

40 Sebastian Moss, “Data Center Power Supply Business Bender Hit by Ransomware Attack,” *Data Center Dynamics*, December 3, 2024, <https://www.datacenterdynamics.com/en/news/data-center-power-supply-business-bender-hit-by-ransomware-attack/>.

41 Scott Robinson, Gavin Wright, and Alexander S. Gillis, “What Is a Side-Channel Attack?,” TechTarget, April 8, 2025, <https://www.techtarget.com/searchsecurity/definition/side-channel-attack>.

42 Gyana Swain, “AMD Discloses New CPU Flaws That Can Enable Data Leaks via Timing Attacks,” CSO, July 10, 2025, <https://www.csoonline.com/article/4020192/amd-discloses-new-cpu-flaws-that-can-enable-data-leaks-via-timing-attacks.html>.

43 Maundrill, “Cybersecurity Implications of Data Centres,” <https://www.infosecurityeurope.com/en-gb/blog/regulation-and-policy/cybersecurity-of-data-centres-uk.html>; “Defining Insider Threats,” Cybersecurity and Infrastructure Security Agency, accessed June 18, 2025, <https://www.cisa.gov/topics/physical-security/insider-threat-mitigation/defining-insider-threats>.

44 Matt Vincent, “How Tariffs Could Impact Data Centers, AI, and Energy amid Supply Chain Shifts,” *Data Center Frontier*, April 3, 2025, <https://www.datacenterfrontier.com/hyperscale/article/55279670/how-tariffs-could-impact-data-centers-ai-and-energy-amid-supply-chain-shifts>; “Securing the Hardware Supply Chain,” OPSWAT, accessed June 18, 2025, <https://www.opswat.com/blog/securing-the-hardware-supply-chain>.

- 45 “Data Center Threats and Vulnerabilities,” Check Point Software Technologies, <https://www.checkpoint.com/cyber-hub/cyber-security/what-is-data-center/data-center-threats-and-vulnerabilities/>.
- 46 “Reimagining Secure Infrastructure for Advanced AI,” OpenAI, May 3, 2024, <https://openai.com/index/reimagining-secure-infrastructure-for-advanced-ai/>.
- 47 Tim Fist, interview by Seungmin Lee, April 24, 2025.
- 48 “GPU Vulnerability: Side-Channel Attacks,” Liquid Web, accessed July 15, 2025, <https://www.liquidweb.com/gpu/vulnerability/#side-channel-attacks>.
- 49 Tim Fist, interview by Seungmin Lee, April 24, 2025.
- 50 “GPU Vulnerability: Side-Channel Attacks,” Liquid Web, <https://www.liquidweb.com/gpu/vulnerability/#side-channel-attacks>.
- 51 Davey Winder, “Nvidia Security Warning—Act Now as 7 New GPU Vulnerabilities Confirmed,” *Forbes*, January 28, 2025, <https://www.forbes.com/sites/daveywinder/2025/01/28/nvidia-security-warning-act-now-as-7-new-gpu-vulnerabilities-confirmed/>.
- 52 Chaim Gartenberg, “TPU Transformation: A Look Back at 10 Years of Our AI-Specialized Chips,” Google Cloud, July 31, 2024, <https://cloud.google.com/transform/ai-specialized-chips-tpu-history-gen-ai>.
- 53 Nate Nelson, “With ‘TPUxtract,’ Attackers Can Steal Orgs’ AI Models,” *DarkReading*, December 13, 2024, <https://www.darkreading.com/vulnerabilities-threats/tpuxtract-attackers-steal-ai-models>.
- 54 “Risks of Data Exfiltration,” SentinelOne, <https://www.sentinelone.com/cybersecurity-101/cybersecurity/data-exfiltration/#risks-of-data-exfiltration>.
- 55 “Case Study: How TPUXtract Leveraged Keysight Tools for AI Model Extraction,” <https://www.keysight.com/blogs/en/tech/nwvs/2025/03/19/case-study-how-tpuxtract-leveraged-keysight-tools-for-ai-model-extraction>.
- 56 Satariano and Mozur, “The Global AI Divide,” <https://www.nytimes.com/interactive/2025/06/23/technology/ai-computing-global-divide.html>.
- 57 “Top 14 AI Security Risks in 2024,” SentinelOne, accessed June 18, 2025, <https://www.sentinelone.com/cybersecurity-101/data-and-ai/ai-security-risks/#14-ai-security-risks-and-threats>.
- 58 “Top 14 AI Security Risks in 2024,” <https://www.sentinelone.com/cybersecurity-101/data-and-ai/ai-security-risks/#14-ai-security-risks-and-threats>.
- 59 Muath Alduhishy, “Sovereign AI: What It Is, and 6 Strategic Pillars for Achieving It,” World Economic Forum, April 25, 2024, <https://www.weforum.org/stories/2024/04/sovereign-ai-what-is-ways-states-building/>.
- 60 Wes Shinego, “Defense Officials Outline AI’s Strategic Role in National Security,” U.S. Department of Defense, April 23, 2025, <https://www.defense.gov/News/News-Stories/Article/Article/4165279/defense-officials-outline-ais-strategic-role-in-national-security/>.
- 61 Shingeo, “Defense Officials Outline AI’s Strategic Role in National Security,” <https://www.defense.gov/News/News-Stories/Article/Article/4165279/defense-officials-outline-ais-strategic-role-in-national-security/>.
- 62 Alduhishy, “Sovereign AI,” <https://www.weforum.org/stories/2024/04/sovereign-ai-what-is-ways-states-building/>.

- 63 Shlomit Wagman, “Weaponized AI: A New Era of Threats and How We Can Counter It,” Harvard Kennedy School Ash Center, April 8, 2025, <https://ash.harvard.edu/articles/weaponized-ai-a-new-era-of-threats/>.
- 64 Carrie Pallardy, “What CISOs Need to Know About Nation-State Actors,” *InformationWeek*, December 12, 2023, <https://www.informationweek.com/cyber-resilience/what-cisos-need-to-know-about-nation-state-actors>.
- 65 Erica D. Lonergan and Michael Poznansky, “A Tale of Two Typhoons: Properly Diagnosing Chinese Cyber Threats,” *War on the Rocks*, February 25, 2025, <https://warontherocks.com/2025/02/a-tale-of-two-typhoons-properly-diagnosing-chinese-cyber-threats/>.
- 66 “Reimagining Secure Infrastructure for Advanced AI,” <https://openai.com/index/reimagining-secure-infrastructure-for-advanced-ai/>.
- 67 Jeremie Harris and Edouard Harris, *America’s Superintelligence Project* (Gladstone AI, April 2025), <https://superintelligence.gladstone.ai/>.
- 68 Billy Perrigo, “Exclusive: Every AI Datacenter Is Vulnerable to Chinese Espionage, Report Says,” *Time*, April 22, 2025, <https://time.com/7279123/ai-datacenter-superintelligence-china-trump-report/>.
- 69 Harris and Harris, *America’s Superintelligence Project*, <https://superintelligence.gladstone.ai>.
- 70 Satariano and Mozur, “The Global AI Divide,” <https://www.nytimes.com/interactive/2025/06/23/technology/ai-computing-global-divide.html>; Zoe Hawkins, Vili Lehdonvirta, and Boxi Wu, “AI Compute Sovereignty: Infrastructure Control Across Territories, Cloud Providers, and Accelerators,” SSRN, June 24, 2025, <https://ssrn.com/abstract=5312977>.
- 71 Amy Gunia, “Will ‘Massive’ Gulf Deals Cement the U.S. Lead in the Race for Global AI Dominance?,” CNN, May 22, 2025, <https://www.cnn.com/2025/05/15/middleeast/trump-abu-dhabi-ai-center-latam-intl>.
- 72 Tye Graham and Peter W. Singer, “How China’s Tech Giants Wired the Gulf,” *Defense One*, May 13, 2025, <https://www.defenseone.com/threats/2025/05/china-tech-giants-wired-gulf/405283/>; Butts, “Malaysia Is Emerging as a Data Center Powerhouse,” <https://www.cnn.com/2024/06/17/malaysia-emerges-as-asian-data-center-powerhouse-amid-booming-demand.html>.
- 73 Google Threat Intelligence Group, “Adversarial Misuse of Generative AI,” Google Cloud, January 29, 2025, <https://cloud.google.com/blog/topics/threat-intelligence/adversarial-misuse-generative-ai>.
- 74 Andersen Cheng, “The Race to Build Data Centers Is On—Here’s How We Keep Them Secure,” *TechRadar Pro*, December 4, 2024, <https://www.techradar.com/pro/the-race-to-build-data-centers-is-on-heres-how-we-keep-them-secure>.
- 75 Marius Hobbhahn, “Scheming Reasoning Evaluations,” Apollo Research, January 23, 2025, <https://www.apolloresearch.ai/research/scheming-reasoning-evaluations>.
- 76 Austin Carson, interview by Seungmin Lee, April 17, 2025.
- 77 Cheng, “The Race to Build Data Centers Is On,” <https://www.techradar.com/pro/the-race-to-build-data-centers-is-on-heres-how-we-keep-them-secure>.
- 78 Tim Fist, interview by Seungmin Lee, April 24, 2025; Fist and Datta, *How to Build the Future of AI in the United States*, <https://ifp.org/future-of-ai-compute>.
- 79 Sella Nevo et al., *Securing AI Model Weights: Preventing Theft and Misuse of Frontier Models* (RAND, 2024), https://www.rand.org/pubs/research_reports/RRA2849-1.html.

80 Nevo et al., *Securing AI Model Weights*, https://www.rand.org/pubs/research_reports/RRA2849-1.html.

81 Fist and Datta, *How to Build the Future of AI in the United States*, <https://ifp.org/future-of-ai-compute>.

82 Vladimir Antić et al., “Protecting Data at Risk of Unintentional Electromagnetic Emanation: TEMPEST Profiling,” *Applied Sciences* 14, no. 11 (June 3, 2024): 4830, <https://doi.org/10.3390/app14114830>.

83 Swati Srivastava, “Regulate or Innovate? Governing AI amid the Race for AI Sovereignty,” *New America*, May 1, 2025, <https://www.newamerica.org/planetary-politics/blog/regulate-or-innovate-governing-ai-amid-the-race-for-ai-sovereignty/>.

84 Arnab Datta and Tim Fist, *Compute in America: A Policy Playbook* (Institute for Progress, February 3, 2025), <https://ifp.org/special-compute-zones/>.

85 Computer Security Resource Center, “Secure Software Development Framework (SSDF),” National Institute of Standards and Technology, updated February 27, 2025, <https://csrc.nist.gov/projects/ssdf>.

86 Ron Ross and Victoria Pillitteri, *Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations*, Special Publication 800-171, rev. 3 (National Institute of Standards and Technology, May 2024), <https://csrc.nist.gov/pubs/sp/800/171/r3/final>.

87 Joint Task Force Working Group, *Security and Privacy Controls for Information Systems and Organizations*, Special Publication 800-53, rev. 5, update 1 (National Institute of Standards and Technology, October 2024), <https://csrc.nist.gov/pubs/sp/800/53/r5/upd1/final>.

88 National Institute of Standards and Technology (NIST), *Security Requirements for Cryptographic Modules*, FIPS PUB 140-3 (NIST, March 22, 2019), <https://doi.org/10.6028/NIST.FIPS.140-3>.

89 “FedRAMP,” Government Services Administration, updated March 31, 2025, <https://www.gsa.gov/technology/government-it-initiatives/fedramp>.

90 U.S. Department of Defense Chief Information Officer, “Cybersecurity Maturity Model Certification,” accessed July 24, 2025, <https://dodcio.defense.gov/CMMC/Model/>.

91 “Zero Trust Maturity Model,” Cybersecurity and Infrastructure Security Agency, accessed April 11, 2023, <https://www.cisa.gov/zero-trust-maturity-model>.

92 “Software Bill of Materials (SBOM),” Cybersecurity and Infrastructure Security Agency, accessed June 21, 2025, <https://www.cisa.gov/sbom>.

93 Biden, *Executive Order on Advancing United States Leadership in Artificial Intelligence*, <https://bidenwhitehouse.archives.gov/briefing-room/presidential-actions/2025/01/14/executive-order-on-advancing-united-states-leadership-in-artificial-intelligence-infrastructure/>.

94 Donald J. Trump, *Executive Order 14179: Removing Barriers to American Leadership in Artificial Intelligence*, 90 FR 874, (The White House, January 31, 2025), <https://www.federalregister.gov/documents/2025/01/31/2025-02172/removing-barriers-to-american-leadership-in-artificial-intelligence>; Secretary of the Interior, *Secretary’s Order No. 3418: Unleashing American Energy* (U.S. Department of the Interior, February 3, 2025), https://www.doi.gov/sites/default/files/document_secretarys_orders/so-3418-signed.pdf; “DOE Identifies 16 Federal Sites Across the Country for Data Center and AI Infrastructure Development,” Department of Energy, April 3, 2025, <https://www.energy.gov/articles/doe-identifies-16-federal-sites-across-country-data-center-and-ai-infrastructure>.

95 *Winning the Race: America's AI Action Plan* (The White House, July 2025), accessed August 7, 2025, <https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf>.

96 *Winning the Race*, <https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf>.

97 Srivastava, "Regulate or Innovate?," <https://www.newamerica.org/planetary-politics/blog/regulate-or-innovate-governing-ai-amid-the-race-for-ai-sovereignty/>.

98 Mariami Tkeshelashvili and Tiffany Saade, *Navigating AI Compliance, Part 2: Risk Mitigation Strategies for Safeguarding Against Future Failures* (Institute for Security and Technology, March 2025), <https://www.securityandtechnology.org/virtual-library/reports/navigating-ai-compliance-part-2/>.

99 "Data Center Threats and Vulnerabilities," Check Point Software Technologies, <https://www.checkpoint.com/cyber-hub/cyber-security/what-is-data-center/data-center-threats-and-vulnerabilities/>; Nevo et al., *Securing Artificial Intelligence Model Weights*, https://www.rand.org/pubs/research_reports/RRA2849-1.html.

100 Fist and Datta, *How to Build the Future of AI in the United States*, <https://ifp.org/future-of-ai-compute>.

101 Cybersecurity and Infrastructure Security Agency (CISA), *Shifting the Balance of Cybersecurity Risk—Principles and Approaches for Secure by Design Software* (CISA, October 25, 2023), <https://www.cisa.gov/resources-tools/resources/secure-by-design>.

102 CISA, *Shifting the Balance of Cybersecurity Risk*, <https://www.cisa.gov/resources-tools/resources/secure-by-design>; "Software Bill of Materials (SBOM)," CISA, <https://www.cisa.gov/sbom>.



This report carries a Creative Commons Attribution 4.0 International license, which permits re-use of New America content when proper attribution is provided. This means you are free to share and adapt New America’s work, or include our content in derivative works, under the following conditions:

- **Attribution.** You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.

For the full legal code of this Creative Commons license, please visit creativecommons.org.

If you have any questions about citing or reusing New America content, please visit www.newamerica.org.

All photos in this report are supplied by, and licensed to, [shutterstock.com](https://www.shutterstock.com) unless otherwise stated. Photos from federal government sources are used under section 105 of the Copyright Act.