

AUGUST 2026

THE INVISIBLE BACKBONE:

How to Secure AI Infrastructure
Today for Tomorrow



DR. TERRENCE KELLY
JESS MAKSIMOV




EXECUTIVE SUMMARY

Artificial intelligence (AI) has increasingly become an underlying infrastructure for the United States, as it is embedded in government operations, national defense, critical infrastructures, and the private sector. Yet the AI sector lacks formal recognition as critical infrastructure, leaving it without public-private coordination mechanisms, information sharing, asset mapping, and protection standards commensurate with AI's increasing role in critical operations.

Simultaneously, the AI sector faces threats that existing federal frameworks were not designed to address. First, adversarial attacks, including data poisoning, distillation, and supply chain manipulation, can exploit the characteristics of AI systems that differ from those of conventional software. Second, because only a handful of providers build the foundation models that other sectors depend on, vulnerabilities in one model can propagate across the codebases built atop it. Third, AI is growing increasingly interdependent with existing critical infrastructure, expanding the attack surface of both the AI sector and the systems it is becoming intertwined with.

To address these risks, the President must designate AI as a critical infrastructure sector. The designation would allow the U.S. government to identify a lead AI infrastructure agency, define which entities constitute the sector, establish industry-wide coordination mechanisms, and develop AI-specific national security standards. To this end, Congress should also empower the executive to set security baselines for AI infrastructure, build federal capacity to coordinate and share threat information, and invest in AI safety and security research and development.





INTRODUCTION

Artificial intelligence (AI) systems are becoming increasingly critical for strengthening U.S. national security. Take Anthropic's Claude Mythos model preview. The system reportedly demonstrates unprecedented AI cyber capabilities, which are key to identifying national security risks. Anthropic claims that Mythos has identified thousands of high- and critical-severity vulnerabilities,¹ a significant jump from just a year ago when AI models were discovering vulnerabilities in single and double digits.² Mythos demonstrates how AI-enhanced attacks could exploit systems while deployers, such as hospitals and utilities, race to patch newly-identified vulnerabilities. Simultaneously, AI is growing more significant for U.S. national security and economic prosperity, making models like Mythos valuable targets for foreign adversaries to distort, replicate, and degrade.³

The stakes extend well beyond any single model, though. The AI systems that already support many U.S. services remain vulnerable to attacks that could take them offline, or worse, allow them to perform seemingly well, while malicious code covertly permeates the infrastructure's backbone.⁴ Because these vulnerabilities cross corporate and infrastructural lines, no single company can defend against them alone. This paper argues that the AI sector must be protected as a whole, through a critical infrastructure designation, increased public-private coordination, information sharing mechanisms, and enhanced infrastructure security research and development (R&D).



¹ Anthropic has professional security contractors currently validating every identified bug report. See Frontier Red Team, "Claude Mythos Preview \ Red.Anthropic.Com," Anthropic, April 7, 2026, <https://red.anthropic.com/2026/mythos-preview/>.

² Christopher Covino, "Policy Memo: Mythos and the Evolving Cyber Landscape," Institute for AI Policy and Strategy, April 16, 2026, https://static1.squarespace.com/static/64edf8e7f2b10d716b5ba0e1/t/69e12e9874e83a6c836164ef/1776365208126/Policy+Memo+_Mythos+and+the+Evolving+Cyber+Landscape.pdf.

³ Covino, "Policy Memo: Mythos and the Evolving Cyber Landscape."

⁴ Junyan Huang et al., "Application of Artificial Intelligence in Medical Imaging for Tumor Diagnosis and Treatment: A Comprehensive Approach," *Discover Oncology* 16 (August 2025): 1625, <https://doi.org/10.1007/s12672-025-03307-3>; Kyle Crichton et al., "Securing Critical Infrastructure in the Age of AI," *Center for Security and Emerging Technology*, October 2024, <https://doi.org/10.51593/20240032>.

DEFINING CRITICAL INFRASTRUCTURE

Before assessing whether AI meets the critical infrastructure threshold, it is necessary to understand what “critical infrastructure” (CI) means and how the United States determines which sectors earn that designation. First, the Critical Infrastructures Protection Act of 2001 defines critical infrastructure as systems and assets “so vital to the United States that the incapacity or destruction of such systems and assets would have a debilitating impact on security, national economic security, national public health or safety, or any combination of those matters.”⁵

Second, as directed by the 2021 National Defense Authorization Act, the Cybersecurity and Infrastructure Security Agency (CISA) proposed a decision-support framework for assessing which sectors should be classified as critical infrastructure.⁶ Under this framework:

01 A COHERENT SECTOR A collection of assets, systems, or networks providing common functions to the economy, government, or society.	02 DEBILITATING IF DISRUPTED Disruptions would produce debilitating impacts on national security, economic security, public health, or safety.
03 A GOVERNANCE GAP Existing regulatory mechanisms, policies, and governance structures do not adequately manage sector-wide risks.	04 COORDINATION PAYS OFF The sector would meaningfully benefit from the coordination and collaboration mechanisms a CI designation offers.

Source: CISA, *FY 2021 National Defense Authorization Act Report*.

Using the CISA framework, this paper defines the AI infrastructure sector, identifies its unique vulnerabilities, maps its interdependencies with other CI sectors, and proposes policy recommendations. This paper serves as the overture in a series that will examine national security threats to different components of the AI infrastructure stack.

⁵ 42 U.S. Code § 5195c - Critical Infrastructures Protection,” Cornell Law School, LII / Legal Information Institute, 2001, <https://www.law.cornell.edu/uscode/text/42/5195c>.

⁶ The 2021 National Defense Authorization Act (NDAA) directed the Secretary of Homeland Security to review existing critical infrastructure sectors and recommend additional designations. Pursuant to the NDAA, CISA was also required to recommend “a process and structure for identifying and designating sectors and aligning those sectors to designated SRMAs [Sector Risk Management Agencies].” *FY 2021 National Defense Authorization Act*, 9002(b) (Cybersecurity and Infrastructure Security Agency, 2021), https://www.nextgov.com/media/section_9002_ndaa_report_final_11.15.21_%281%29.pdf.

⁷ CISA, *FY 2021 National Defense Authorization Act Report*.

What is the AI Sector?

This paper defines the AI sector from an infrastructure lens, as comprising organizations, facilities, technologies, and industries whose primary purpose is the development, training, deployment, and operation of AI systems. This sector includes the hardware, software, and operational infrastructure that directly perform AI computation or produce the specialized components required for it.

The AI sector includes four primary components: brains, brawn, building blocks, and benefit delivery. The first three facilitate AI production and services, and the last delivers these systems to customers.

01 BRAINS AI MODEL DEVELOPMENT & SOFTWARE Model design, training, model weights, evaluation and alignment systems, and specialized software — developed by frontier labs and research organizations.	02 BRAWN AI COMPUTE INFRASTRUCTURE Specialized facilities, such as data centers, and the hardware that runs AI workloads.
03 BUILDING BLOCKS AI HARDWARE DESIGN & MANUFACTURING Specialized chip design, advanced semiconductor production in fabrication plants, and the chip supply chain.	04 BENEFIT DELIVERY AI OPERATIONAL & DELIVERY SYSTEMS The platforms and infrastructure that deploy and serve AI models at scale, translating raw capabilities into usable services.

The first three components facilitate AI production and services; the fourth delivers them to customers. Source: ARI.

The AI sector also depends on external critical infrastructure, such as energy, communications, information technology (IT), finance, water and wastewater, and transportation. These remain distinct sectors when assessing interdependencies and vulnerabilities. However, for the purposes of this paper, outside infrastructure becomes part of the AI sector when it is purpose-built to perform AI computational functions or when it becomes an inherent part of said functions during operations. For example, the communications infrastructure that connects AI data centers forms the transit part of the AI sector,⁸ while transmission and on-site generation campuses constitute the electricity portion of these facilities.

⁸ Note that most CIs have interdependencies with other CIs, as is discussed in the sections below.

The AI Backbone of Modern Services

Government Services and Defense

AI systems are increasingly integrated into CI functions, forming interdependencies as these sectors rely on AI and AI relies on them. Between August 2022 and August 2023, AI-related federal contracts increased by almost 150 percent to \$675 million.⁹ U.S. government services use AI systems for enterprise-wide functions and individual productivity. Federal agencies already deploy AI tools to carry out their mandates. The Food and Drug Administration (FDA) uses AI systems to review drug applications, the Centers for Disease Control and Prevention (CDC) uses them to analyze medical images for abnormalities, the Department of (DOT) uses machine learning AI to predict flight delays, and U.S. Customs and Border Protection (CBP) deploys AI systems to analyze border crossing patterns.¹⁰ A 2024 Ernst & Young survey also found that 51 percent of state and local public sector employees, and 64 percent of federal employees, use AI applications several times a week or daily.¹¹

FIGURE 01 DAILY AI USE IN GOVERNMENT

Share of U.S. government employees who say they use AI daily



Source: Ernst & Young, *Insights into the Integration of AI in Government*, 2024 pulse survey of U.S. government employees.

The U.S. military already relies on AI systems for predictive maintenance and military logistics, and is increasingly turning to AI technologies for military planning and operations.¹² In a January 2026 memorandum, Secretary of War Pete Hegseth instructed the Department of War (DOW) to implement AI tools with speed, noting that “the risks of not moving fast enough outweigh the risks of imperfect alignment.”¹³

⁹ Sella Nevo, “A Sprint Toward Security Level 5,” IFP, August 11, 2025, <https://ifp.org/a-sprint-toward-security-level-5/>; Kevin C. Desouza, “The Evolution of Artificial Intelligence (AI) Spending by the U.S. Government,” Brookings, March 26, 2024, <https://www.brookings.edu/articles/the-evolution-of-artificial-intelligence-ai-spending-by-the-u-s-government/>.

¹⁰ Sanam Hooshidary et al., “Artificial Intelligence in Government: The Federal and State Landscape,” National Conference of State Legislatures, November 22, 2024, <https://www.ncsl.org/technology-and-communication/artificial-intelligence-in-government-the-federal-and-state-landscape>.

¹¹ Amy Jones, “Insights into the Integration of AI in Government,” Ernst & Young Global Limited, September 18, 2024, https://www.ey.com/en_us/industries/government-public-sector/insights-into-the-integration-of-ai-in-government.

¹² Sanam Hooshidary et al., “Artificial Intelligence in Government: The Federal and State Landscape.”

¹³ Pete Hegseth, “Artificial Intelligence Strategy for the Department of War,” *Memorandum for Senior Pentagon Leadership Commanders of the Combatant Commands Defense Agency and DOW Field Activity Directors*, January 9, 2026, <https://media.defense.gov/2026/Jan/12/2003855671/-1/-1/0/artificial-intelligence-strategy-for-the-department-of-war.pdf>.

U.S. military AI adoption is increasing. In 2026, Deputy Secretary of War Steve Feinberg designated Palantir's Maven Smart System (MSS) as a program of record to allow for the AI system's eventual "enterprise-wide integration" into the U.S. military's command and control systems.¹⁴ The U.S. military has already deployed AI tools on classified data to provide real-time targeting in the U.S. military operations in Iran.¹⁵ In addition, open source reporting indicates that AI tools were used in U.S. operations to capture the former Venezuelan president, although the exact use cases remain classified.¹⁶ The Central Intelligence Agency (CIA) is also reportedly deploying AI systems to help analyze the plans, intentions, and capabilities of foreign nations, with CIA Deputy Director Michael Ellis revealing that the agency recently used AI to create the CIA's first autonomous intelligence report.¹⁷

The Private Sector

AI integration is also growing across private CI sectors.¹⁸ Financial institutions, for instance, are adopting AI systems for fraud detection, forecasting, trading, lending, cybersecurity, and customer service.¹⁹ A 2024 Treasury Department report notes that these institutions have relied on AI-based fraud detection for over a decade and increasingly apply AI to cybersecurity and code generation.²⁰ But deployment carries risks: in 2025, the Financial Stability Oversight Council warned that "widespread [AI] adoption" in finance must be monitored so that agencies and institutions can manage threats from "malicious actors, as well as national security, cyber, or other unanticipated risks."²¹

The energy sector is following suit, integrating AI into predictive grid maintenance, anomalous event detection, and energy management and efficiency technologies to support grid reliability.²² According to a 2024 Department of Energy (DOE) report, researchers and industry actors have long used machine learning models to analyze grid data drawn from sensors, meters, and power plants.²³ The advent of generative AI now extends these use cases to grid planning, permitting and siting, grid operations and resilience, and security. However, the DOE researchers warn that foundation models can produce stochastic outputs without mitigation measures for energy applications.

¹⁴ Brandi Vincent, "Feinberg's New Maven Directive Sets AI-Enabled Decision-Making as 'the Cornerstone' for CJADC2," *DefenseScoop*, April 3, 2026, <https://defensescoop.com/2026/04/03/palantir-maven-feinberg-directive/>.

¹⁵ Tara Copp et al., "Anthropic's AI Tool Claude Central to U.S. Campaign in Iran, amid a Bitter Feud," *The Washington Post*, March 4, 2026, <https://www.washingtonpost.com/technology/2026/03/04/anthropic-ai-iran-campaign/>.

¹⁶ Amrith Ramkumar et al., "Exclusive: Pentagon Used Anthropic's Claude in Maduro Venezuela Raid," *Tech, Wall Street Journal*, February 13, 2026, <https://www.wsj.com/politics/national-security/pentagon-used-anthropics-claude-in-maduro-venezuela-raid-583aff77>.

¹⁷ John Sakellariadis, "CIA Is Trusting AI to Help Analyze Intel from Human Spies," *PoliticoPro*, April 9, 2026, <https://subscriber.politicopro.com/article/2026/04/cia-ai-intelligence-analysis-00865893>.

¹⁸ According to an MIT study, as of 2025, 80 percent of surveyed corporations had explored individual generative AI productivity tools, such as Copilot and ChatGPT, with 40 percent deploying these; 60 percent of the surveyed organizations evaluated enterprise-grade generative AI tools, with five percent reaching production. See Aditya Challapally et al., *The GenAI Divide: State of AI in Business 2025*, July 2025, https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf; Alexandra Sasha Luccioni et al., "From Efficiency Gains to Rebound Effects: The Problem of Jevons' Paradox in AI's Polarized Environmental Debate," *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, June 23, 2025, 76–88, <https://doi.org/10.1145/3715275.3732007>.

¹⁹ In a 2025 McKinsey & Company survey of 120 chief financial officers, respondents reported using generative AI for decision support (retrieving data, generating reports, and running forecasts or scenarios), capital management (improving efficiency in payable and receivable processes), and cost optimization (applying advanced algorithms to spot anomalies and areas of waste in invoice and purchase order data). See Alexander Sukharevsky et al., "AI in Finance: Driving Automation and Business Value," McKinsey & Company, November 3, 2025, <https://www.mckinsey.com/capabilities/strategy-and-corporate-finance/our-insights/how-finance-teams-are-putting-ai-to-work-today>; U.S. Department of the Treasury, *Artificial Intelligence Report on the Uses, Opportunities, and Risks of Artificial Intelligence in the Financial Services Sector U.S. Department of the Treasury December 2024 in Financial Services* (2024), <https://home.treasury.gov/system/files/136/Artificial-Intelligence-in-Financial-Services.pdf>; 2025 FSOC Annual Report (Financial Stability Oversight Council, 2025), <https://home.treasury.gov/system/files/261/FSOC2025AnnualReport.pdf>; Ken Sweet, "Mastercard Expects to Find Compromised Cards Quicker Using AI," May 22, 2024, <https://apnews.com/article/mastercard-visa-ai-credit-card-fraud-detection-0c348818087a57b13bfac66c761e03b4>.

²⁰ Keith Benes et al., *AI for Energy: Opportunities for a Modern Grid and Clean Energy Economy*, April 2024, https://www.energy.gov/sites/default/files/2024-04/AI%20EO%20Report%20Section%205.2gfi_043024.pdf; U.S. Department of the Treasury, *Managing Artificial Intelligence Specific Cybersecurity Risks In The Financial Services Sector* (2024), <https://www.northerntrust.com/content/dam/northerntrust/events/documents/us-dept-of-treasury-report-managing-ai.pdf>; "Treasury Announces Enhanced Fraud Detection Processes, Including Machine Learning AI, Prevented and Recovered Over \$4 Billion in Fiscal Year 2024," U.S. Department of the Treasury, October 17, 2024, <https://home.treasury.gov/news/press-releases/jy2650>.

²¹ 2025 FSOC Annual Report.

²² Shane Londagin, "Unlocking AI's Grid Modernization Potential," *Federation of American Scientists*, June 25, 2025, <https://fas.org/publication/unlocking-ai-grid-modernization-potential/>.


²³ Benes et al., *AI for Energy: Opportunities for a Modern Grid and Clean Energy Economy*.

In the IT sector, AI systems already write code and support cybersecurity systems.²⁴ First, AI-assisted coding now underpins the software and platforms on which U.S. sectors increasingly rely. A 2025 survey found that 51 percent of professional code developers use AI tools in their development processes daily.²⁵ Microsoft's Chief Executive Officer, Satya Nadella, claims that AI now writes approximately 30 percent of the company's code, and GitHub reports that its Copilot AI tool writes an average of 46 percent of code across all programming languages.²⁶

FIGURE 02 SECURITY INCIDENTS FROM AI-GENERATED CODE

"Has your organization ever identified a security vulnerability introduced by AI-generated code?"

1 IN 5 SUFFERED A SERIOUS INCIDENT



Source: Aikido, *State of AI in Security & Development*, 2026.

Second, AI cybersecurity systems can identify malicious events or threats; detect anomalies; and, with generative AI, proactively mitigate code vulnerabilities and analyze threat actor behavior across various sectors.²⁷ AI-supplemented cybersecurity use cases date back to the 1990s, when they scanned security logs for anomalies.²⁸ Now, with the emergence of generative AI systems, a PYMNTS report found that between May 2023 and August 2024, the percentage of chief operating officers adopting AI-powered cybersecurity management systems within their companies increased from 17 to 55 percent.²⁹ Moreover, Anthropic's Mythos model is reportedly unearthing decade-old vulnerabilities and writing exploits at an accelerated pace that exceeds traditional patching timelines.³⁰

All of these developments introduce new risks. A 2026 Aikido study found that 20 percent of surveyed organizations experienced serious security vulnerabilities due to AI-generated code.³¹ And, in July 2025, an attacker exploited a flaw in the Amazon Q Developer (a coding assistant tool that developers install in their coding environments), injecting malicious instructions into the official product, which was distributed through a widely used marketplace. The compromised extension had 964,000 installations.³² Ultimately, the code did not execute only because of a syntax error in its instructions.³³

²⁴ Madeline Field, "Your Defense Code Is Already AI-Generated. Now What?," *War on the Rocks*, March 25, 2026, <https://warontherocks.com/cogs-of-war/your-defense-code-is-already-ai-generated-now-what/>; Crichton et al., "Securing Critical Infrastructure in the Age of AI."

²⁵ "AI: 2025 Stack Overflow Developer Survey, 2025 Developer Survey, 2025, <https://survey.stackoverflow.co/2025/ai#sentiment-and-usage-ai-select-ai-sel-prof>.

²⁶ Note that these companies have market incentives to provide upper estimates of their AI use. See Maxwell Zeff, "Microsoft CEO Says up to 30% of the Company's Code Was Written by AI," *TechCrunch*, April 30, 2025, <https://techcrunch.com/2025/04/29/microsoft-ceo-says-up-to-30-of-the-companys-code-was-written-by-ai/>; Thomas Dohmke, "GitHub Copilot for Business Is Now Available," *The GitHub Blog*, February 14, 2023, <https://github.blog/news-insights/product-news/github-copilot-for-business-is-now-available/>.

²⁷ Crichton et al., "Securing Critical Infrastructure in the Age of AI."

²⁸ Ibid.

²⁹ PYMNTS Intelligence "COOs Leverage GenAI to Reduce Data Security Losses," *PYMNTS.com*, December 2024, https://www.pymnts.com/study_posts/coos-leverage-genai-to-reduce-data-security-losses/.

³⁰ Robert Wiblin, "How Scary Is Claude Mythos? 303 Pages in 21 Minutes," *80,000 Hours*, April 10, 2026, <https://80000hours.org/2026/04/claude-mythos-hacking-alignment/>.

³¹ Aikido, "State of AI in Security & Development," accessed April 21, 2026, <https://www.aikido.dev/state-of-ai-security-development-2026-download>; Field, "Your Defense Code Is Already AI-Generated. Now What?," *War on the Rocks*, March 25, 2026, <https://warontherocks.com/cogs-of-war/your-defense-code-is-already-ai-generated-now-what/>.

³² Ibid.

³³ Amazon Web Services, "Security Update for Amazon Q Developer Extension for Visual Studio Code (Version #1.84)," AWS, July 23, 2025, <https://aws.amazon.com/security/security-bulletins/AWS-2025-015/>; <https://aws.amazon.com/security/security-bulletins/AWS-2025-015/>.

Unique Threats & Vulnerabilities

The same integration that makes AI valuable also makes it a target, introducing new attack vectors and vulnerabilities that are distinct from traditional cyber threats. Specifically, generative AI outputs are context-specific and nondeterministic; therefore, an AI model's degradation is difficult to identify purely based on how it functions.³⁴ To secure these AI systems end-to-end, developers must maintain AI safety (ensuring robust systems are developed through the data collection, training, and deployment phases) and AI security (guarding against external threats).³⁵

How might a threat actor target the AI sector? First, actors can poison the system's training data or use prompt injection attacks to cause unintended model behaviors, disrupt the system's effectiveness, or extract sensitive information.³⁶ Similar to traditional cyberattacks, adversaries can exploit AI systems by implanting malicious code. However, due to AI systems' generative nature and pervasiveness in code writing, this manipulation may go undetected if model outputs and the software they support seemingly function correctly.³⁷

Second, attackers can exfiltrate model weights, which store AI systems' core intelligence and cost billions of dollars in data, compute, and algorithms to develop.³⁸ Compromising these weights would give attackers access to an AI developer's most valuable assets, enabling adversaries to exploit the technology or use it to develop competing models.³⁹ Similarly, attackers can use "distillation" techniques to prompt the model's inference application programming interface (API) with repeated queries until they expose the model's functionality; they can then use this information to train new models that mimic the original's behavior.⁴⁰ In April 2026, the Office of Science and Technology Policy released National Security and Technology Memorandum 4, which identified "deliberate, industrial-scale campaigns to distill U.S. frontier AI systems" and directed federal agencies to engage with the private sector on information sharing and mitigation measures.⁴¹

Third, actors can leverage global AI manufacturing and supply chain chokepoints to exploit third-party suppliers of U.S. AI hardware.⁴² States can use targeted export controls, import restrictions, or sanctions on critical minerals (e.g., China's restrictions on gallium and germanium) to create semiconductor delivery delays and shortages. Actors can also ship semiconductor-grade materials with minor impurities, causing chip fabrication yields to plummet, or sabotage concentrated manufacturing clusters, such as those in East Asia, to disrupt hardware supply chains.⁴³

Fourth, using side-channel attacks on AI data centers, actors can physically or remotely track emissions, including acoustic and power fluctuations, to infer proprietary and security secrets such as encryption keys.⁴⁴ According to a 2024 RAND report, building the secure data centers and hardware needed to defend against such sophisticated attacks requires sector-wide, national-security-oriented R&D that remains underdeveloped.⁴⁵

³⁴ A 2024 RAND study identified 38 "meaningfully distinct attack vectors" for frontier model weight theft. See Sella Nevo et al., *Securing AI Model Weights: Preventing Theft and Misuse of Frontier Models* (RAND, 2024), https://www.rand.org/pubs/research_reports/RR_A2849-1.html; Jeremie Harris and Edouard Harris, "America's Superintelligence Project," America's Superintelligence Project, April 2025, <https://superintelligence.gladstone.ai/>; Sella Nevo, "A Sprint Toward Security Level 5," Institute for Progress, Emerging Technology, *Emerging Technology*, August 11, 2025, <https://ifp.org/a-sprint-toward-security-level-5/>; Field, "Your Defense Code Is Already AI-Generated. Now What?," March 25, 2026; Erich Grunewald and Asher Brass, "Accelerating AI Data Center Security," Institute for AI Policy and Strategy, September 9, 2025, <https://www.iaps.ai/research/accelerating-ai-data-center-security>.

³⁵ Crichton et al., "Securing Critical Infrastructure in the Age of AI."

³⁶ Ibid.

³⁷ Nevo et al., *Securing AI Model Weights*.

³⁸ Ibid.

³⁹ Ibid.

⁴⁰ Ibid.

⁴¹ Office of Science and Technology Policy, "NSTM-4: Memorandum for the Heads of Executive Departments and Agencies," April 23, 2026, <https://www.whitehouse.gov/wp-content/uploads/2026/04/NSTM-4.pdf>.

⁴² Nevo et al., *Securing AI Model Weights*.

⁴³ Terrence Kelly, "AI Advanced Manufacturing" (unpublished manuscript, 2026).

⁴⁴ Grunewald and Brass, "Accelerating AI Data Center Security."

⁴⁵ Nevo et al., *Securing AI Model Weights*.

Moreover, AI systems introduce automation bias, which can undermine a user's ability to determine whether AI system outputs are reliable and secure. A review of 74 studies across healthcare, aviation, and military sectors shows that when an automated decision support system provides an output, human oversight degrades. Therefore, AI safety and security are vital to ensuring that the systems humans increasingly rely on and trust are indeed reliable.⁴⁶

Under Threat: Why AI Is Already a Target

Assessing a target's risk requires weighing three factors: threat, vulnerability, and consequence. To be considered a threat, an actor must have the intent and capability to produce harm,⁴⁷ be able to identify a sector's exploitable vulnerabilities, and inflict consequences on the selected target. Since threat actors would not be able to disrupt all U.S. critical infrastructure systems simultaneously and have finite resources, they would have to set priorities. To determine targets, adversaries often consider which infrastructures would have the greatest consequences, are the easiest to attack (most vulnerable), and bring the lowest escalation risks.

AI tools run on foundation models that are concentrated under a handful of companies with substantially overlapping training data. The National Institute of Standards and Technology (NIST) notes that in this "algorithmic monoculture," many consequential decision-making sectors use the same model or algorithm, which can result in increased susceptibility to correlated failures.⁴⁸ Therefore, compromising a foundation model can propagate vulnerabilities across many of the systems built atop it.⁴⁹

There are already documented cases of adversaries targeting and using AI infrastructure. In March 2026, Iran attacked Amazon Web Services (AWS) data centers in the United Arab Emirates, striking U.S. cloud and AI chokepoints, with the state media describing the operation as strikes on "the enemy's technological infrastructure."⁵⁰ In addition, state-linked Shenzhen University accessed Nvidia A100 and H100 chips via AWS, revealing their susceptibility to remote access.⁵¹

Interdependency: Vulnerabilities & Disruptions

Due to the AI sector's dependence on critical infrastructures and its concentration in a handful of models, the disruption of AI functionality through interconnected sectors, such as energy and communications, can have cascading impacts on U.S. services.

AI & Energy

The Energy Paper in this series identifies transmission as the most vulnerable link for AI functionality in the energy sector. First, U.S. transmission is constrained by inadequate capacity and decade-long timelines for permitting and building new high-voltage lines.⁵² Second, the U.S. grid is disaggregated across the country. This structure limits power sharing, meaning localized generation failures cannot be easily offset by drawing power

⁴⁶ Kate Goddard et al., "Automation Bias: A Systematic Review of Frequency, Effect Mediators, and Mitigators," *Journal of the American Medical Informatics Association* 19, no. 1 (2012), <https://doi.org/https://doi.org/10.1136/amiajnl-2011-000089>.

⁴⁷ Ciaran Martin, "Typhoons in Cyberspace," Royal United Services Institute for Defence and Security Studies, June 22, 2026, <https://www.rusi.org>; Cybersecurity & Infrastructure Security Agency, "PRC State-Sponsored Actors Compromise and Maintain Persistent Access to U.S. Critical Infrastructure," February 7, 2024, <https://www.cisa.gov/news-events/cybersecurity-advisories/aa24-038a>.

⁴⁸ National Institute of Standards and Technology (US), *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*, (National Institute of Standards and Technology (U.S.), 2024), 600-1, <https://doi.org/10.6028/NIST.AI.600-1>.

⁴⁹ Field, "Your Defense Code Is Already AI-Generated. Now What?," March 25, 2026.

⁵⁰ Oliver Jabbour, "When Data Centres Become Targets: It's Time to Start Treating AI Infrastructure as Critical Infrastructure," World Economic Forum, April 2, 2026, <https://www.weforum.org/stories/2026/04/ai-infrastructure-critical-infrastructure/>.

⁵¹ Eduardo Baptista et al., "Chinese Entities Turn to Amazon Cloud and Its Rivals to Access High-End US Chips, AI," Technology, *Reuters*, August 22, 2024, <https://www.reuters.com/technology/chinese-entities-turn-amazon-cloud-its-rivals-access-high-end-us-chips-ai-2024-08-23/>.

⁵² Jason Beckfield et al., *How Grid Projects Get Stuck: Four Cases in Long-Distance Transmission Development in the United States* (Harvard University, 2024), <https://salatainstitute.harvard.edu/wp-content/uploads/2024/06/How-Grid-Projects-Get-Stuck-Report.pdf>.

from other regions. Third, the specialized substations and large power transformers (LPTs) that serve AI data centers are chokepoints; these structures are typically lightly guarded and vulnerable to both physical sabotage and cyberattacks. A successful strike on one of these could take a facility offline for months. Moreover, most AI data centers only maintain short-term backup power, which cannot support long-term disruptions during prolonged attacks.

Separately, grid modernization through smarter control and behind-the-meter generation introduces new cyber vulnerabilities, as these technologies depend on the integrity of grid sensor data, making them susceptible to attacks that can corrupt or interrupt the data flow.⁵³ Disrupting power control systems, particularly those used by grid operators to balance supply and demand, would give an adversary significant leverage. Based on previous state-sponsored infiltration operations by China-linked Salt and Volt Typhoon groups, major adversaries may already possess the capability to cause such disruptions.⁵⁴

Finally, the concentration of large AI data center loads in select regional clusters creates systemic risks. For example, if the facilities' protective controls disconnect unexpectedly during a grid disturbance, the sudden loss of massive loads can destabilize entire regional power systems. This means that a single point of failure could cascade into widespread AI service disruption.⁵⁵

AI & Communications

AI systems are equally dependent on communications infrastructure. First, unlike standard internet applications, AI clusters operate as synchronized supercomputers using specialized, lossless protocols, so even a 0.1 percent packet loss can stall training runs.⁵⁶ Sophisticated adversaries could exploit this sensitivity by inducing latency that degrades performance in time-critical applications, such as financial systems or military operations.

Second, despite the geographic distribution of data centers,⁵⁷ their traffic aggregates through a small number of internet exchange points (IXPs) and carrier hotels, creating significant chokepoints.⁵⁸ A single attack on a high-density peering hub can simultaneously isolate multiple AI inference nodes from their users.⁵⁹ Third, fiber, electricity transmission, and transportation infrastructures frequently share physical corridors,⁶⁰ meaning one event, such as a storm, construction accident, or deliberate attack, can sever these infrastructures simultaneously.⁶¹ Even if a data center is running on backup power, it cannot deliver AI services if its external fiber links are cut.⁶²

⁵³ Note that data centers are not the only virtual power plants on a power network; other entities, from residential solar systems through large industrial systems, can contribute to this effect.

⁵⁴ Rebecca Smith, "Russian Hackers Reach U.S. Utility Control Rooms, Homeland Security Officials Say," *Wall Street Journal*, July 23, 2018, <https://www.wsj.com/articles/russian-hackers-reach-u-s-utility-control-rooms-homeland-security-officials-say-1532388110>; Cybersecurity & Infrastructure Security Agency, "Russian Government Cyber Activity Targeting Energy and Other Critical Infrastructure Sectors," March 16, 2018, <https://www.cisa.gov/news-events/alerts/2018/03/15/russian-government-cyber-activity-targeting-energy-and-other-critical-infrastructure-sectors>.

⁵⁵ Elinor Ginzburg-Ganz et al., "Technical Challenges of AI Data Center Integration into Power Grids—A Survey," *Energies* 19, no. 1 (2026): 137, <https://doi.org/10.3390/en19010137>.

⁵⁶ "Why 0.1% Packet Loss Is Killing Your AI: High-Stakes Lessons from the NVIDIA NCP-AIN," *FlashGenius*, January 2026, <https://flashgenius.net/blog-article/why-01-packet-loss-is-killing-your-ai-high-stakes-lessons-from-the-nvidia-ncp-ain>.

⁵⁷ Eric Geller, "AI's 'Connective Tissue' Is Woefully Insecure, Cisco Warns," *Cybersecurity Dive*, February 19, 2026, <https://www.cybersecuritydive.com/news/ai-agents-model-context-protocol-cisco-report/812580/>.

⁵⁸ Diana Hernández, "Geographic Concentration of Data Centers and Their Future: Resilience of Current Options and New Possibilities," BITA, December 19, 2025, <https://bita.ghost.io/geographic-concentration-of-data-centers-and-their-future-resilience-of-current-options-and-new-possibilities/>.

⁵⁹ Joe Wargo, "Physical Vulnerability of Fiber Compared to Wireless Backhaul," *AOI Wireless*, July 31, 2024, <https://www.aowireless.com/post/physical-vulnerability-of-fiber-compared-to-wireless-backhaul>.

⁶⁰ Asad Ramzanali and Benjamin Dinovelli, *Dig Once: How Federal, State, and Local Governments Can Reduce the Cost of Broadband Deployment* (Vanderbilt Policy Accelerator, 2025).

⁶¹ Joint Committee on the National Security Strategy, *Subsea Telecommunications Cables: Resilience and Crisis Preparedness*, First Report of Session 2024–26 HC 723 / HL Paper 179 (UK Parliament, 2025), <https://publications.parliament.uk/pa/jt5901/jtselect/jtmatsec/723/report.html>.

⁶² Geller, "AI's 'Connective Tissue' Is Woefully Insecure, Cisco Warns."

How Could the Energy and Communications Sectors Be Interdicted?

A threat actor targeting U.S. AI infrastructure would prioritize the systems whose disruption would produce the greatest effect relative to the actor's conflict objectives. Attacks could focus on the grid and communications nodes that are most heavily loaded, hardest to replace, and most critical to key AI functions. In a conflict involving maritime operations, for example, an adversary would likely target the U.S. AI systems that support logistics, ports, and troop movement, rather than striking indiscriminately.

To disrupt power, adversaries' priority targets would include (1) natural gas compressors serving major data center hubs, (2) specialized substations with LPTs,⁶³ and (3) the software platforms managing grid performance. To disrupt communications functions, IXPs and carrier hotels represent the highest-value chokepoints. These attacks could be kinetic or cyber. For example, substations outside guarded facilities are physically vulnerable; in some cases, a single attacker with a weapon can damage them.⁶⁴ Adversaries do not need to instigate full blackouts. Targeted, temporary disruptions to the grid, communications infrastructure, or behind-the-meter systems could degrade AI functionality sufficiently to send strategic signals or restrict U.S. AI services in a conflict.

Space and Bioeconomy: Critical, But Not CI

The practical weight of CISA's CI designation framework becomes clearest when applied to sectors that have sought, and failed to receive, a designation. In 2023, the Cyberspace Solarium Commission 2.0—a think tank created from the congressionally established U.S. national cybersecurity strategy body—recommended that space systems be designated as critical infrastructure because (1) they have been targeted in adversarial attacks and espionage campaigns, (2) other critical infrastructure sectors depend on the technologies, (3) they exhibit unique physical and technical characteristics, and (4) much of U.S. terrestrial infrastructure is located abroad. Similarly, CISA proposed that the bioeconomy and space systems sectors be evaluated under the framework above. However, the Solarium report admits that federal rules already regulate space systems, with the DOW's jurisdiction over defense applications and the Federal Communications Commission's authority over space-based communications systems.⁶⁵

Ultimately, the President did not designate these sectors as CI in National Security Memorandum (NSM)-22, which renewed the U.S. critical infrastructure sector framework in 2024. Senior officials noted that this was due to the space sector already being so ingrained into “many different sectors” that designating it as a standalone entity did not “make sense,” and because risks to the bioeconomy “were not unique enough to merit a new sector.”⁶⁶

While these sectors were ultimately passed over due to existing regulations and lead agencies with jurisdictions over the systems' core functions, no comparable authority currently governs the AI sector's risks at the scale its integration demands. And whereas the bioeconomy's hazards were deemed too similar to those addressed elsewhere to justify a standalone designation, the AI sector confronts vulnerabilities without parallels in other sectors—training-data poisoning, the theft of model weights, distillation of proprietary systems, and a concentration of foundation models in which a single compromise can propagate across everything built upon it. The AI sector thus satisfies the very conditions that space and the bioeconomy could not: it lacks an existing regulatory regime equal to its risks, and its threats are distinct enough to warrant a CI designation of its own.

⁶³ LPTs are particularly attractive targets because they can take months or even years to replace.

⁶⁴ See, for example, Evan Halper and Marc Lifsher, “Attack on California Electric Grid Called ‘Terrorism,’” *Government Technology*, February 10, 2014, <https://www.govtech.com/em/disaster/attack-electric-grid-raises-alarm-em.html>.

⁶⁵ Frank Cilluffo et al., *Time to Designate Space Systems as Critical Infrastructure* (Cyberspace Solarium Commission 2.0, 2023); Christian Vasquez, “Cyberspace Solarium Commission Says Space Systems Should Be Considered Critical Infrastructure,” *Cyberscoop*, April 14, 2023, <https://cyberscoop.com/solarium-commission-space-systems-critical-infrastructure/>.

⁶⁶ Mariam Baksh, “CISA Highlights Space, Bioeconomy as Possible New Critical Infrastructure Sectors,” *Nextgov/FCW*, November 16, 2022, <https://www.nextgov.com/cybersecurity/2022/11/cisa-highlights-space-bioeconomy-possible-new-critical-infrastructure-sectors/379818/>; Eric Geller, “The White House Has a New Master Plan to Stop Worst-Case Scenarios,” *WIRED*, April 30, 2024, <https://www.wired.com/story/biden-national-security-memorandum-critical-infrastructure-threats/>; *FY 2021 National Defense Authorization Act*.

Recommendations

Despite AI systems confronting unique threats, private companies lack the resources and multi-sectoral coordination capacity necessary to secure their technologies against nation-state adversaries. The frontier AI labs have commercial incentives to protect their intellectual property and users' privacy against competitors, criminal hacking groups, insider threats, and industrial espionage, but that protection stops at the firm's perimeter.⁶⁷ U.S. government capacity can fill this gap, protecting AI infrastructure at a scale no single company can achieve alone. To secure U.S. AI infrastructure, Americans for Responsible Innovation recommends the following actions:

Designate AI as a Critical Infrastructure Sector

The President must designate the AI sector as critical infrastructure in a national security memorandum, subjecting it to incident reporting requirements under the Cyber Incident Reporting for Critical Infrastructure Act of 2022 (CIRCIA). The memorandum would secure the sector by:⁶⁸

01	Assigning an AI-specific Sector Risk Management Agency (SRMA) to coordinate across the federal government and with the AI industry to help manage and mitigate AI-specific risks while building U.S. governmental capacity; ⁶⁹
02	Incorporating AI into the biennial National Risk Management Plan that summarizes U.S. risk mitigation efforts to the President; ⁷⁰
03	Empowering the SRMA to work across agencies and with the private sector to create security and resilience standards across the entire relevant AI infrastructure stack;
04	Facilitating public-private engagements across interdependent CI sectors; ⁷¹
05	Identifying AI-specific Systemically Important Entities (SIEs) to prioritize government resources, risk mitigation, and security standards;
06	Directing the intelligence community to conduct intelligence estimates on AI sector risks and, as permitted by law and regulation, share those with industry; and
07	Allowing the SRMA to coordinate authorities (e.g., the Defense Production Act), financing, and other capabilities that would enable executive intervention to secure essential supply chains and expedite infrastructure buildouts (e.g., through streamlined permitting). ⁷²

⁶⁷ Nevo et al., *Securing AI Model Weights*.

⁶⁸ Brian E. Humphreys, The 2024 National Security Memorandum on Critical Infrastructure Security and Resilience, In Focus no. IF12716 (Congressional Research Service, 2024), <https://www.congress.gov/crs-product/IF12716>; LII Leg. Inf. Inst., "42 U.S. Code § 5195c - Critical Infrastructures Protection"; Cybersecurity & Infrastructure Security Agency, "Critical Infrastructure Sectors," Cisa.Gov, accessed June 22, 2026, <https://www.cisa.gov/topics/critical-infrastructure-security-and-resilience/critical-infrastructure-sectors>; Michael Gruden, "Changes to Critical Infrastructure Requirements," Crowell & Moring, January 28, 2025, <https://www.crowell.com/en/insights/publications/changes-to-critical-infrastructure-requirements>; Justine Phillips, "Cyber Incident Reporting for Critical Infrastructure Act of 2022 (CIRCIA)," Cybersecurity Fact Sheet, 2024; Consolidated Appropriations Act, 2022, H.R.2471, 117th Congress (2021-2022) (2022), <https://www.congress.gov/bills/117/congress/house-bill/2471/text>.

⁶⁹ Megan Brown et al., "New White House Policy Previews Increased Cybersecurity Oversight and Regulation," Wiley Rein LLP, May 2, 2024, <https://www.wiley.law/alert-New-White-House-Policy-Previews-Increased-Cybersecurity-Oversight-and-Regulation>.

⁷⁰ "National Security Memorandum on Critical Infrastructure Security and Resilience," CISA, accessed April 20, 2026, <https://www.cisa.gov/national-security-memorandum-critical-infrastructure-security-and-resilience>.

⁷¹ Humphreys, The 2024 National Security Memorandum on Critical Infrastructure Security and Resilience.

⁷² Arnab Datta and Tim Fist, "Compute in America: A Policy Playbook," Institute for Progress, Emerging Technology, February 3, 2025, <https://ifp.org/special-compute-zones/>; Asad Ramzanali, "How to Regulate the Cloud: A Blueprint to Address the Market Failures and National Security Risks of Cloud Computing," SSRN Scholarly Paper no. 6222822 (Social Science Research Network, September 18, 2025), <https://doi.org/10.2139/ssrn.6222822>.

In addition, CISA and NIST should conduct a risk-based assessment to classify systems according to their frontier level, intended use cases, deployed sector, and potential impact on human health, safety, and security.⁷³ This would create a dynamic federal framework for selecting the entities that constitute the AI sector and mitigating their vulnerabilities.

Establish Security Baselines

To mitigate AI data center security vulnerabilities, CISA and the Center for AI Standards and Innovation (CAISI) should publish updated AI data center guidance establishing security baselines.⁷⁴ To achieve this, Congress can direct CISA to create a task force to audit the physical, cyber, and personnel security of AI data centers; identify vulnerabilities; and work with AI companies to develop enforceable security standards across the AI infrastructure stack.

Spur AI Infrastructure Security R&D

In parallel, to build a nascent AI infrastructure safety and security R&D ecosystem, Congress should allocate funding to the Defense Advanced Research Projects Agency (DARPA) to administer grants that create a cost-effective market for the technology required to achieve AI infrastructure standards calibrated to national security threats.⁷⁵

Taken together, these measures would give the federal government the coordination capacity that the AI threat environment already demands, and that the private sector alone cannot supply.

Conclusion

The AI sector already bears the hallmarks of critical infrastructure. It is interwoven with public and private services, concentrated among a handful of foundation models, and increasingly interdependent with CI sectors, meaning a single attack on the AI stack could cascade across multiple sectors at once.

The U.S. threat environment compounds these structural concerns. State-linked actors have already struck AI data centers and demonstrated sophisticated penetration capabilities, suggesting that the conditions for disruption are taking shape. Simultaneously, strategic assets, such as Mythos and similar successors, are precisely the targets that adversaries will be incentivized to attack. Every month of delay widens the gap between how much the nation relies on AI and how little it protects it. As these threats grow, the President and Congress must act to fortify U.S. AI infrastructure and expand public-private coordination in this burgeoning sector.

⁷³ Bipartisan Policy Center, “Defining High-Risk, High-Reward AI,” Explainer, April 6, 2023, <https://bipartisanpolicy.org/explainer/high-risk-high-reward-ai/>; NIST, AI Risk Management Framework, NIST AI 100-1 (National Institute of Standards and Technology, 2023), <https://doi.org/10.6028/NIST.AI.100-1>; European Commission, “AI Act: Shaping Europe’s Digital Future,” June 19, 2026, <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>; “Annex III: High-Risk AI Systems Referred to in Article 6(2),” EU Artificial Intelligence Act, n.d., accessed June 22, 2026, <https://artificialintelligenceact.eu/annex/3/>.

⁷⁴ “Center for AI Standards and Innovation (CAISI),” NIST, October 26, 2023, <https://www.nist.gov/caisi>; “AI Data Security: Best Practices for Securing Data Used to Train & Operate AI Systems,” Joint Cybersecurity Information, May 2025, https://media.defense.gov/2025/May/22/2003720601/-1/-1/0/CSI_AI_DATA_SECURITY.PDF; Nevo, “A Sprint Toward Security Level 5.”

⁷⁵ “R&D Opportunities,” DARPA, Defense Advanced Research Projects Agency, accessed June 22, 2026, <https://www.darpa.mil/work-with-us/opportunities>; Grunewald and Brass, “Accelerating AI Data Center Security.”

Bibliography

- “Annex III: High-Risk AI Systems Referred to in Article 6(2).” EU Artificial Intelligence Act. Accessed June 22, 2026. <https://artificialintelligenceact.eu/annex/3/>.
- Aikido, “State of AI in Security & Development.” 2026. <https://www.aikido.dev/state-of-ai-security-development> 2026-download.
- Amazon Web Services. “Security Update for Amazon Q Developer Extension for Visual Studio Code (Version #1.84).” AWS, Inc., July 23, 2025. <https://aws.amazon.com/security/security-bulletins/AWS-2025-015/>.
- Ansolabehere, Stephen, Parrish Bergquist, Major Eason, Daniel Alain Evrard, Quinn Lewis, Elizabeth Thom, Jason Beckfield, and Ana Martinez. How Grid Projects Get Stuck: Four Cases in Long-Distance Transmission Development in the United States. Harvard University, 2024. <https://salatainstitute.harvard.edu/wp-content/uploads/2024/06/How-Grid-Projects-Get-Stuck-Report.pdf>.
- Arora, Chhavi, Marc Sorel, Arjita Bhan, Jess He, Nicholas Shaw, Pankaj Sachdeva, Riya Garg, and Shriya Ravishankar. “The next Big Shifts in AI Workloads and Hyperscaler Strategies.” McKinsey & Company, December 17, 2025. <https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/the-next-big-shifts-in-ai-workloads-and-hyperscaler-strategies>?
- Baksh, Mariam. “CISA Highlights Space, Bioeconomy as Possible New Critical Infrastructure Sectors.” Nextgov/FCW, November 16, 2022. <https://www.nextgov.com/cybersecurity/2022/11/cisa-highlights-space-bioeconomy-possible-new-critical-infrastructure-sectors/379818/>.
- Baptista, Eduardo, Fanny Potkin, and Karen Freifeld. “Chinese Entities Turn to Amazon Cloud and Its Rivals to Access High-End US Chips, AI.” Technology. Reuters, August 22, 2024. <https://www.reuters.com/technology/chinese-entities-turn-amazon-cloud-its-rivals-access-high-end-us-chips-ai-2024-08-23/>.
- Benes, Keith, Joshua Porterfield, and Charles Yang. AI for Energy: Opportunities for a Modern Grid and Clean Energy Economy. U.S. Department of Energy, April 2024. [https://www.energy.gov/sites/default/files/2024-04/AI%20EO%20Report%20Section%205.2g\(i\)_043024.pdf](https://www.energy.gov/sites/default/files/2024-04/AI%20EO%20Report%20Section%205.2g(i)_043024.pdf).
- Bick, Alexander, Adam Blandin, and David Deming. “The State of Generative AI Adoption in 2025.” Federal Reserve Bank of St. Louis, November 13, 2025. <https://www.stlouisfed.org/on-the-economy/2025/nov/state-generative-ai-adoption-2025>.
- Biden, Joseph R. “National Security Memorandum on Critical Infrastructure Security and Resilience.” The American Presidency Project, April 30, 2024. <https://www.presidency.ucsb.edu/documents/national-security-memorandum-critical-infrastructure-security-and-resilience>.
- Bipartisan Policy Center. “Defining High-Risk, High-Reward AI.” Explainer, April 6, 2023. <https://bipartisanpolicy.org/explainer/high-risk-high-reward-ai/>.
- Brown, Megan, Kathleen Scott, Jaclyn Brown, Sydney White, and Joshua Waldman. “New White House Policy Previews Increased Cybersecurity Oversight and Regulation.” Wiley Rein LLP, May 2, 2024. <https://www.wiley.law/alert-New-White-House-Policy-Previews-Increased-Cybersecurity-Oversight-and-Regulation>.
- Buchanan, Bonnie G. Artificial Intelligence in Finance. The Alan Turing Institute, April 2019.
- Carlini, Nicholas, Keane Lucas, Evyatar Ben Asher, Newton Cheng, Hasnain Lakhani, David Forsythe, and Kyla Guru. “0-Days \ Red.Anthropic. Com.” February 5, 2026. <https://red.anthropic.com/2026/zero-days/>.
- Challapally, Aditya, Chris Pease, Ramesh Raskar, and Pradyumna Chari. The GenAI Divide: State of AI in Business 2025. July 2025. https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf.
- Cilluffo, Frank, Mark Montgomery, Sharon Cardash, and Kelsey Shields. Time to Designate Space Systems as Critical Infrastructure. Cyberspace Solarium Commission 2.0, 2023. https://cybersolarium.org/wp-content/uploads/2023/04/CSC2.0_Report_SpaceCriticalInfrastructureSector.pdf.
- Copp, Tara, Elizabeth Dwoskin, and Ian Duncan. “Anthropic’s AI Tool Claude Central to U.S. Campaign in Iran, Amid a Bitter Feud.” The Washington Post, March 4, 2026. <https://www.washingtonpost.com/technology/2026/03/04/anthropic-ai-iran-campaign/>.
- Covino, Christopher. “Policy Memo: Mythos and the Evolving Cyber Landscape.” Institute for AI Policy and Strategy, April 16, 2026. https://static1.squarespace.com/static/64edf8e7f2b10d716b5ba0e1/t/69e12e9874e83a6c836164ef/1776365208126/Policy+Memo_+Mythos+and+the+Evolving+Cyber+Landscape.pdf.
- Crichton, Kyle, Jessica Ji, Kyle Miller, John Bansemmer, Zachary Arnold, David Batz, Minwoo Choi, Marisa Decillis, Patricia Eke, Daniel M. Gerstein, Alex Leblang, Monty McGee, Greg Rattray, Luke Richards, and Alana Scott. “Securing Critical Infrastructure in the Age of AI.” Center for Security and Emerging Technology, October 2024. <https://doi.org/10.51593/20240032>.

- Cybersecurity & Infrastructure Security Agency. “Critical Infrastructure Sectors.” CISA. Accessed June 22, 2026. [https:// www.cisa.gov/topics/critical-infrastructure-security-and-resilience/critical-infrastructure-sectors](https://www.cisa.gov/topics/critical-infrastructure-security-and-resilience/critical-infrastructure-sectors).
- Cybersecurity & Infrastructure Security Agency. “PRC State-Sponsored Actors Compromise and Maintain Persistent Access to U.S. Critical Infrastructure.” CISA. February 7, 2024. <https://www.cisa.gov/news-events/cybersecurity-advisories/aa24-038a>.
- Cybersecurity & Infrastructure Security Agency. “Russian Government Cyber Activity Targeting Energy and Other Critical Infrastructure Sectors.” March 16, 2018. <https://www.cisa.gov/news-events/alerts/2018/03/15/russian-government-cyber-activity-targeting-energy-and-other-critical-infrastructure-sectors>.
- Cybersecurity and Infrastructure Security Agency, “National Security Memorandum on Critical Infrastructure Security and Resilience.” April 30, 2024. [https:// www.cisa.gov/national-security-memorandum-critical-infrastructure-security-and-resilience](https://www.cisa.gov/national-security-memorandum-critical-infrastructure-security-and-resilience).
- Cybersecurity and Infrastructure Security Agency, FY 2021 National Defense Authorization Act. 9002(b). CISA, 2021. [https://www. nextgov.com/media/section_9002_ndaa_report_final_11.15.21_%281%29.pdf](https://www.nextgov.com/media/section_9002_ndaa_report_final_11.15.21_%281%29.pdf).
- Cybersecurity and Infrastructure Security Agency. “Cyber Incident Reporting for Critical Infrastructure Act of 2022 (CIR CIA): Fact Sheet.” CISA, 2022. <https://www.cisa.gov/resources-tools/resources/cyber-incident-reporting-critical-infrastructure-act-2022-fact-sheet>.
- Datta, Arnab and Tim Fist. “Compute in America: A Policy Playbook.” Institute for Progress. Emerging Technology. Emerging Technology, February 3, 2025. <https://ifp.org/special-compute-zones/>.
- Defense Advanced Research Projects Agency. “R&D Opportunities.” DARPA. Accessed June 22, 2026. [https://www.darpa.mil/ work-with-us/opportunities](https://www.darpa.mil/work-with-us/opportunities).
- Dohmke, Thomas. “GitHub Copilot for Business Is Now Available.” The GitHub Blog, February 14, 2023. [https://github.blog/ news-insights/product-news/github-copilot-for-business-is-now-available/](https://github.blog/news-insights/product-news/github-copilot-for-business-is-now-available/).
- El Hajj, Mohammad and Jamil Hammoud. “Unveiling the Influence of Artificial Intelligence and Machine Learning on Financial Markets: A Comprehensive Analysis of AI Applications in Trading, Risk Management, and Financial Operations.” Journal of Risk and Financial Management 16, no. 10 (2023): 434. <https://doi.org/10.3390/jrfm16100434>.
- European Commission. “AI Act: Shaping Europe’s Digital Future.” June 19, 2026. [https://digital-strategy.ec.europa.eu/en/ policies/regulatory-framework-ai](https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai).
- Field, Madeline. “Your Defense Code Is Already AI-Generated. Now What?” War on the Rocks, March 25, 2026. [https:// warontherocks.com/cogs-of-war/your-defense-code-is-already-ai-generated-now-what/](https://warontherocks.com/cogs-of-war/your-defense-code-is-already-ai-generated-now-what/).
- Financial Stability Oversight Council. 2025 FSOC Annual Report. 2025. [https://home.treasury.gov/system/files/261/ FSOC2025AnnualReport.pdf](https://home.treasury.gov/system/files/261/FSOC2025AnnualReport.pdf).
- Flash Genius, “Why 0.1% Packet Loss Is Killing Your AI: High-Stakes Lessons from the NVIDIA NCP-AIN.” FlashGenius, January 2026. [https:// flashgenius.net/blog-article/why-01-packet-loss-is-killing-your-ai-high-stakes-lessons-from-the-nvidia-ncp-ain](https://flashgenius.net/blog-article/why-01-packet-loss-is-killing-your-ai-high-stakes-lessons-from-the-nvidia-ncp-ain).
- Frontier Red Team. “Claude Mythos Preview \ Red.Anthropic.Com.” Anthropic, April 7, 2026. [https://red.anthropic.com/2026/ mythos-preview/](https://red.anthropic.com/2026/mythos-preview/).
- Geller, Eric. “AI’s ‘Connective Tissue’ Is Woefully Insecure, Cisco Warns.” Cybersecurity Dive, February 19, 2026. [https://www. cybersecuritydive.com/news/ai-agents-model-context-protocol-cisco-report/812580/](https://www.cybersecuritydive.com/news/ai-agents-model-context-protocol-cisco-report/812580/).
- Geller, Eric. “The White House Has a New Master Plan to Stop Worst-Case Scenarios.” WIRED, April 30, 2024. <https://www.wired.com/story/biden-national-security-memorandum-critical-infrastructure-threats/>.
- Ginzburg-Ganz, Elinor, Pavel Lifshits, Ram Machlev, Juri Belikov, Ziv Krieger, and Yoash Levron. “Technical Challenges of AI Data Center Integration into Power Grids—A Survey.” Energies 19, no. 1 (2026): 137. <https://doi.org/10.3390/en19010137>.
- Goddard, Kate, Abudl Roudsari, and Jeremy Wyatt. “Automation Bias: A Systematic Review of Frequency, Effect Mediators, and Mitigators.” Journal of the American Medical Informatics Association 19, no. 1 (2012). <https://doi.org/10.1136/amiajnl-2011-000089>.
- Gruden, Michael. “Changes to Critical Infrastructure Requirements.” Crowell & Moring, January 28, 2025. [https://www. crowell.com/en/insights/publications/changes-to-critical-infrastructure-requirements](https://www.crowell.com/en/insights/publications/changes-to-critical-infrastructure-requirements).

- Grunewald, Erich and Asher Brass. “Accelerating AI Data Center Security.” Institute for AI Policy and Strategy, Institute for AI Policy and Strategy, September 9, 2025. <https://www.iaps.ai/research/accelerating-ai-data-center-security>.
- Halper, Evan and Marc Lifsher. “Attack on California Electric Grid Called ‘Terrorism.’” GovTech, February 10, 2014. <https://www.govtech.com/em/disaster/Attack-Electric-Grid-Raises-Alarm-EM.html>.
- Harris, Jeremie and Edouard Harris. “America’s Superintelligence Project.” America’s Superintelligence Project, April 2025. <https://superintelligence.gladstone.ai/>.
- Hegseth, Pete. “Artificial Intelligence Strategy for the Department of War: Memorandum for Senior Pentagon Leadership Commanders of the Combatant Commands Defense Agency and DOW Field Activity Directors.” Department of War. January 9, 2026. <https://media.defense.gov/2026/Jan/12/2003855671/-1/-1/0/artificial-intelligence-strategy-for-the-department-of-war.pdf>.
- Hernández, Diana. “Geographic Concentration of Data Centers and Their Future: Resilience of Current Options and New Possibilities.” BITA, December 19, 2025. <https://bita.ghost.io/geographic-concentration-of-data-centers-and-their-future-resilience-of-current-options-and-new-possibilities/>.
- Hooshidary, Sanam, Chelsea Canada, and William Clark. “Artificial Intelligence in Government: The Federal and State Landscape.” National Conference of State Legislatures, November 22, 2024. <https://www.ncsl.org/technology-and-communication/artificial-intelligence-in-government-the-federal-and-state-landscape>.
<https://doi.org/10.5281/ZENODO.2612537>.
- Huang, Junyan, Yizhen Xiang, Shengqi Gan, Linrong Wu, Jiangyu Yan, Dong Ye, and Junjun Zhang. “Application of Artificial Intelligence in Medical Imaging for Tumor Diagnosis and Treatment: A Comprehensive Approach.” Discover Oncology 16 (August 2025): 1625. <https://doi.org/10.1007/s12672-025-03307-3>.
- Humphreys, Brian E. The 2024 National Security Memorandum on Critical Infrastructure Security and Resilience. In Focus No. IF12716. Congressional Research Service, 2024. <https://www.congress.gov/crs-product/IF12716>.
- Hunter, Benjamin, Sumeet Hindocha, and Richard W. Lee. “The Role of Artificial Intelligence in Early Cancer Diagnosis.” Cancers 14, no. 6 (2022): 1524. <https://doi.org/10.3390/cancers14061524>.
- Jabbour, Oliver. “When Data Centres Become Targets: It’s Time to Start Treating AI Infrastructure as Critical Infrastructure.” World Economic Forum, April 2, 2026. <https://www.weforum.org/stories/2026/04/ai-infrastructure-critical-infrastructure/>.
- Jeffries, Hakeem S. Consolidated Appropriations Act, 2022, H.R.2471, 117th Congress (2021-2022). <https://www.congress.gov/bill/117th/congress/house-bill/2471/text>.
- Joint Committee on the National Security Strategy. Subsea Telecommunications Cables: Resilience and Crisis Preparedness. First Report of Session 2024–26 HC 723 / HL Paper 179. UK Parliament, 2025. <https://publications.parliament.uk/pa/jt5901/jtselect/jtnatsec/723/report.html>.
- Jones, Amy. “Insights into the Integration of AI in Government.” Ernst & Young Global Limited, September 18, 2024. https://www.ey.com/en_us/industries/government-public-sector/insights-into-the-integration-of-ai-in-government.
- Kostova, Tamara. “GenAI Vs. Robo-Advisors: Considerations For The Financial Industry.” Forbes. Accessed April 21, 2026. <https://www.forbes.com/councils/forbesbusinesscouncil/2024/08/15/genai-vs-robo-advisors-considerations-for-the-financial-industry/>.
- Larson, Jacob, James S. Denford, Gregory S. Dawson, and Kevin C. Desouza. “The Evolution of Artificial Intelligence (AI) Spending by the U.S. Government.” Brookings, March 26, 2024. <https://www.brookings.edu/articles/the-evolution-of-artificial-intelligence-ai-spending-by-the-u-s-government/>.
- Legal Information Institute. “42 U.S. Code § 5195c - Critical Infrastructures Protection.” Cornell Law School. 2001. <https://www.law.cornell.edu/uscode/text/42/5195c>.
- Londagin, Shane. “Unlocking AI’s Grid Modernization Potential.” Federation of American Scientists, June 25, 2025. <https://fas.org/publication/unlocking-ai-grid-modernization-potential/>.
- Luccioni, Alexandra Sasha, Emma Strubell, and Kate Crawford. “From Efficiency Gains to Rebound Effects: The Problem of Jevons’ Paradox in AI’s Polarized Environmental Debate.” Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency, June 23, 2025, 76–88. <https://doi.org/10.1145/3715275.3732007>.

Ma, Jiwon. "Time to Designate Space Systems as Critical Infrastructure." Cyberspace Solarium Commission 2.0, April 14, 2023. <https://cybersolarium.org/csc-2-0-reports/time-to-designate-space-systems-as-critical-infrastructure/>.

Martin, Ciaran. "Typhoons in Cyberspace." Royal United Services Institute for Defence and Security Studies, June 22, 2026. <https://www.rusi.orghttps://www.rusi.org>.

National Institute of Standards and Technology, "Center for AI Standards and Innovation (CAISI)." NIST,

National Institute of Standards and Technology. AI Risk Management Framework. NIST AI 100-1. NIST, 2023. <https://doi.org/10.6028/NIST.AI.100-1>.

National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. NIST, 2024. <https://doi.org/10.6028/NIST.AI.600-1>.

National Security Agency's Artificial Intelligence Security Center, Cybersecurity and Infrastructure Security Agency, Federal Bureau of Investigation, Australian Cyber Security Centre, New Zealand National Cyber Security Centre, and United Kingdom National Cyber Security Centres. "AI Data Security: Best Practices for Securing Data Used to Train & Operate AI Systems." Joint Cybersecurity Information, May 2025. https://media.defense.gov/2025/May/22/2003720601/-1/-1/0/CSI_AI_DATA_SECURITY.PDF.

Nevo, Sella, Dan Lahav, Ajay Karpur, Yogev Bar-On, Henry Alexander Bradley, and Jeff Alstott. Securing AI Model Weights: Preventing Theft and Misuse of Frontier Models. RAND, 2024. https://www.rand.org/pubs/research_reports/RR2849-1. html.

Nevo, Sella. "A Sprint Toward Security Level 5." Institute for Progress. August 11, 2025. <https://ifp.org/a-sprint-toward-security-level-5/>.

October 26, 2023. <https://www.nist.gov/caisi>.

Office of Science and Technology Policy. "NSTM-4: Memorandum for the Heads of Executive Departments and Agencies." April 23, 2026. <https://www.whitehouse.gov/wp-content/uploads/2026/04/NSTM-4.pdf>.

Organization for Economic Co-operation and Development. Competition in the Provision of Cloud Computing Services. OECD Competition Law and Policy Working Papers. 323rd ed. OECD Competition Law and Policy Working Papers. 2025. <https://doi.org/10.1787/595859c5-en>.

PYMNTS Intelligence. "COOs Leverage GenAI to Reduce Data Security Losses." PYMNTS, December 2024. https://www.pymnts.com/study_posts/coos-leverage-genai-to-reduce-data-security-losses/.

Ramkumar, Amrith, Keach Hagey, and Vera Bergengruen. "Exclusive: Pentagon Used Anthropic's Claude in Maduro Venezuela Raid." Tech. Wall Street Journal, February 13, 2026. <https://www.wsj.com/politics/national-security/pentagon-used-anthropics-claude-in-maduro-venezuela-raid-583aff17>.

Ramzanali, Asad and Benjamin Dinovelli. Dig Once: How Federal, State, and Local Governments Can Reduce the Cost of Broadband Deployment. Vanderbilt Policy Accelerator, 2025. <https://cdn.vanderbilt.edu/vu-URL/wp-content/uploads/sites/412/2025/12/12003228/Dig-Once.pdf>.

Ramzanali, Asad. "How to Regulate the Cloud: A Blueprint to Address the Market Failures and National Security Risks of Cloud Computing." SSRN Scholarly Paper No. 6222822. Social Science Research Network, September 18, 2025. <https://doi.org/10.2139/ssrn.6222822>.

Rashid, Adib Bin and Ashfakul Karim Kausik. "AI Revolutionizing Industries Worldwide: A Comprehensive Overview of Its Diverse Applications." Hybrid Advances 7 (December 2024): 100277. <https://doi.org/10.1016/j.hybadv.2024.100277>.

Sakellariadis, John. "CIA Is Trusting AI to Help Analyze Intel from Human Spies." PoliticoPro, April 9, 2026. <https://subscriber.politicopro.com/article/2026/04/cia-ai-intelligence-analysis-00865893>.

Smith, Rebecca. "Russian Hackers Reach U.S. Utility Control Rooms, Homeland Security Officials Say." Wall Street Journal, July 23, 2018. <https://www.wsj.com/articles/russian-hackers-reach-u-s-utility-control-rooms-homeland-security-officials-say-1532388110>.

Stack Overflow. "AI: 2025 Stack Overflow Developer Survey." 2025. <https://survey.stackoverflow.co/2025/> ai#sentiment-and-usage-ai-select-ai-sel-prof.

Sukharevsky, Alexander, Andy West, Cristina Catania, and David Grande. "AI in Finance: Driving Automation and Business Value." McKinsey & Company, November 3, 2025. <https://www.mckinsey.com/capabilities/strategy-and-corporate-finance/our-insights/how-finance-teams-are-putting-ai-to-work-today>.

Sweet, Ken. "Mastercard Expects to Find Compromised Cards Quicker Using AI." AP News, May 22, 2024. <https://apnews.com/article/mastercard-visa-ai-credit-card-fraud-detection-0c348818087a57b13bfac66c761e03b4>.

- U.S. Department of the Treasury. “Treasury Announces Enhanced Fraud Detection Processes, Including Machine Learning AI, Prevented and Recovered Over \$4 Billion in Fiscal Year 2024.” October 17, 2024. <https://home.treasury.gov/news/press-releases/jy2650>.
- U.S. Department of the Treasury. “Treasury Releases Report on the Uses, Opportunities, and Risks of Artificial Intelligence in Financial Services.” February 13, 2026. <https://home.treasury.gov/news/press-releases/jy2760>.
- U.S. Department of the Treasury. Artificial Intelligence Report on the Uses, Opportunities, and Risks of Artificial Intelligence in the Financial Services Sector. 2024. <https://home.treasury.gov/system/files/136/Artificial-Intelligence-in-Financial-Services.pdf>.
- U.S. Department of the Treasury. Managing Artificial Intelligence Specific Cybersecurity Risks In the Financial Services Sector. 2024. <https://www.northerntrust.com/content/dam/northerntrust/events/documents/us-dept-of-treasury-report-managing-ai.pdf>.
- Vasquez, Christian. “Cyberspace Solarium Commission Says Space Systems Should Be Considered Critical Infrastructure.” Cyberscoop, April 14, 2023. <https://cyberscoop.com/solarium-commission-space-systems-critical-infrastructure/>.
- Vincent, Brandi. “Feinberg’s New Maven Directive Sets AI-Enabled Decision-Making as ‘the Cornerstone’ for CJADC2.” DefenseScoop, April 3, 2026. <https://defensescoop.com/2026/04/03/palantir-maven-feinberg-directive/>.
- Wargo, Joe. “Physical Vulnerability of Fiber Compared to Wireless Backhaul.” AOWireless, July 31, 2024. <https://www.aowireless.com/post/physical-vulnerability-of-fiber-compared-to-wireless-backhaul>.
- Wiblin, Robert. “How Scary Is Claude Mythos? 303 Pages in 21 Minutes.” 80,000 Hours, April 10, 2026. <https://80000hours.org/2026/04/claude-mythos-hacking-alignment/>.
- Yan, Sheno, Shen Wang, Yue Duan, Hanbin Hong, Kiho Lee, Doowon Kim, and Yuan Hong. “An LLM-Assisted Easy-to-Trigger Backdoor Attack on Code Completion Models: Injecting Disguised Vulnerabilities against Strong Detection.” arXiv:2406.06822. Preprint, arXiv, June 10, 2024. <https://doi.org/10.48550/arXiv.2406.06822>.
- Zeff, Maxwell. “Microsoft CEO Says up to 30% of the Company’s Code Was Written by AI.” TechCrunch, April 30, 2025. <https://techcrunch.com/2025/04/29/microsoft-ceo-says-up-to-30-of-the-companys-code-was-written-by-ai/>.