MERCATUS CENTER
George Mason University

POLICY BRIEF

Critical Risks: Rethinking Critical Infrastructure Policy for Targeted AI Regulation

Matthew Mittelsteadt

March 2024

As Congress and the Biden administration move into 2024, the boundless artificial intelligence (AI) policy conversation is shifting toward certain lowest common denominators. Perhaps chief among them: interest in action to assure AI safety in critical infrastructure (CI). In recent months, both congressional chambers have put special emphasis on AI's impact on CI. In December, the House Homeland Security Committee held a hearing on AI safety in critical infrastructure,[1] while in the Senate, a recent "AI Insight Forum" discussed possible regulatory action.[2] While most such policy messaging is just talk, AI is a rare area where hints of real action are emerging. Recently, a sizable bipartisan collection of senators led by Senators John Thune (R-SD) and Amy Klobuchar (D-MN) introduced the bipartisan Artificial Intelligence Research, Innovation, and Accountability Act of 2023, an attempt at light-touch regulation focusing on just the most critical systems.[3] Meanwhile, in the White House's recent AI executive order, CI was given top billing as a major concern. The White House's words have since been translated to action through required AI CI risk assessments, diplomatic efforts to develop guidelines for AI system security in critical systems and infrastructure, and even potential regulatory rules.[4] Although any attempt to go big on AI is inherently difficult, this piece of the AI regulatory pie appears to be picking up modest momentum.

While those regulatory and policy design choices hold ample room for debate, the issue of whether to take action does not. As the Colonial Pipeline hack of 2021 illustrated in sharp relief, our critical infrastructure's deep security issues are already clear: failure carries broad consequences, and any AI-specific risks may indeed warrant consideration of regulation.

Despite any potential justification, however, our current CI policy framework is simply unfit for any AI regulatory task. Implemented in both statute and policy, today's CI policy is a rickety construction whose unclear boundaries, vast scope, and uneven bureaucratic and industry com-

mitment already fail to meet the demands of current nonregulatory policy goals. Under today's policy directives, nearly anything and everything is or can be classified as critical infrastructure, yet seemingly that policy reality remains misunderstood to many, both those demanding action and those actively crafting legislation. Any moves to build regulation on top of this shaky foundation invite unnecessary policy risks, including an unintentionally vast regulatory scope, a lack of prioritization, and regulatory swirl.

AI is still new, and decisions made in the critical few years following ChatGPT's release will carry long-run importance. Rather than rush to regulate, we must work to get this right. Building AI policy on top of this creaky legal and policy frame risks a failure to address real system safety challenges while threatening continued innovation and economic freedom. To proceed, let's first examine (1) why AI CI action is under consideration, (2) what CI comprises, (3) what weaknesses exist in the current CI policy framework, and finally (4) how we can set AI CI policy up for success.

1. Emerging AI CI Concerns

While critical infrastructure policy has been around since the Clinton administration, only in recent years has Congress embraced the need for more ambitious action, even regulation. Specifically, such conversations have been prompted by a worrying string of infrastructure cyberattacks in the 2020s, including the following:

1. **Oldsmar water treatment attack.** In February 2021, cyber intruders compromised the systems of a water treatment plant outside Tampa, Florida, altering sodium hydroxide levels in the water from 100 parts per million to a toxic 11,100 parts per million. Through that simple cyber mischief, trusted drinking water was quickly transformed into caustic poison. Luckily, the problem was detected before mass harm or casualties.[5]
2. **Colonial Pipeline hack.** In May 2021, DarkSide, a Russia-based hacker group, locked the billing systems of the Colonial Pipeline with ransomware. The result was near-instant economic strife. Gas lines piled up, administrators scrambled, and prices surged.[6]
3. **Danish power grid attack.** A large-scale cyberattack launched by the Russian GRU in May 2023 compromised 22 Danish grid operators, forcing companies to disconnect from the grid, thereby destabilizing grid stability.[7]

While each incident speaks to technical insecurity, the Danish attack speaks to an evolving threat landscape. Today, geopolitical tensions are prompting worries about further, more destabilizing state-sponsored security crises. In December, the Iranian Islamic Revolutionary Guard Corps, motivated by the ongoing Israel-Hamas war, attacked a series of American water and wastewater plants.[8] Perhaps more concerning, in January, Christopher Wray, the director of the Federal Bureau of Investigation, reported that the United States had disrupted a large-scale Chinese state-sponsored operation that implanted malware aimed at shutting down targets including water, transportation, and

energy facilities.[9] That incident was unprecedented and represents a significant strategic shift on China's part away from cyber spying, theft, and vandalism toward deadly cyber-physical attacks.

Those incidents show there is reason for fear. The security and safety of the digital systems that operate our infrastructure are already flawed, insecure, and increasingly subject to potentially devastating attack. When failures do occur, consequences can be immense. In February 2021, a series of major ice storms hammered Texas, causing widespread instability in the state's independent power grid.[10] While no resident was without power for longer than three days, the crisis still cost 246 lives by an official estimate, with many unofficial counts estimating a death toll up to four times that amount.[11] No matter the root cause, be it weather, cyberattack, or AI malfunction, systems such as the grid carry little room for error, and failures can put hundreds of lives in danger.

Naturally, these noted risks are *cyber* risks (and weather risks in the case of Texas), not AI-specific problems. Today, AI is not widely used in many of our most critical systems, including water, power, or pipelines; however, applications are quickly maturing. Surveying specifically grid infrastructure, the International Energy Agency notes boundless AI infrastructure uses, including solar and wind weather prediction, distributed device management for increasingly complex grids, demand response balancing, and predictive maintenance, among others.[12] While many applications are just turning the corner, some are proven. In 2019, Google augmented its windfarms with weather prediction models that enabled wind power output predictions 36 hours ahead of time.[13] The result was a dramatic 20 percent increase in revenue per megawatt-hour. That is just one example, yet it illustrates the potentially immense economic and environmental incentives to integrate AI into these systems.

With these improvements, however, there is good reason to question whether certain AI systems could not only assume but also exacerbate already-proven cyber risks. Fundamentally, AI cybersecurity is a dark frontier. According to a machine learning security initiative of the Defense Advanced Research Project Agency, the frontier technology research arm of the Department of Defense, today "a comprehensive theoretical understanding of [machine learning] vulnerabilities is lacking."[14] Because highly capable AI is so new, security researchers simply haven't had the time to understand cyber risks. What research has been done, however, suggests reason for worry. A 2024 report from the National Institute of Standards and Technology illustrates a variety of emerging AI-specific security issues.[15] For instance, AI systems are subject to data-poisoning attacks, whereby a hacker injects vulnerabilities into AI systems by spoiling their upstream training data. System security no longer depends just on airtight code, but on airtight data as well. Another AI-specific insecurity is the transferability of vulnerabilities. As general-purpose models and off-the-shelf code are fine-tuned for bespoke purposes, weaknesses can be transmitted from one system to another, raising the possibility of systemic insecurities.

When it comes to patching these security holes, further challenges emerge. A researcher at the Center for Security and Emerging Technology at Georgetown University notes the uniquely dif-

ficult and costly task of AI risk mitigation.[16] Today, AI vulnerabilities often cannot be identified until after the considerably expensive training process; meanwhile, patching vulnerabilities often requires retraining, costing both time and money. For CI applications with high security demands, these considerable constraints will uniquely stress robust cyber defense.

At present, such security and safety concerns should certainly raise eyebrows, and the use of AI in CI indeed demands due consideration. How to respond to these threats through policy, regulation, or security assistance, however, is a matter of debate, and there is no clear one-size-fits-all policy fix. Furthermore, it's unclear if AI systems are or will be less secure than traditional technologies in coming years. Policymakers must consider whether a bespoke AI security treatment makes sense compared with across-the-board technology-agnostic security efforts. The point here is not to prescribe a solution but to highlight the scale of a potential threat and emphasize why government action, regulatory or not, is likely and even warranted in this case.

No matter the chosen path, however, our current system is simply not set up for success. Solving these problems requires some groundwork.

2. A Critical Look at What Is "Critical"

Today, the biggest CI policy challenge—essential to any targeted regulation—is identifying what exactly "critical" means.

Looking at the words of policy influencers and decision-makers in Washington, we see broad, commonsense agreement on what exactly this term should entail. In December's AI CI hearing, Representative Carlos Giménez (R-TX) described CI as "our electric grid, . . . our piping, . . . and the things that are vital to our everyday life."[17] Largely agreeing with the congressman, the Government Accountability Office, the congressional oversight research agency, recently stated that CI comprises the systems that provide "the essential functions—such as supplying water, generating energy, and producing food—that underpin American society."[18] Finally, even the White House agreed that "the infrastructure that underpins our economy, public health and safety, and national security" means "our power grids, pipelines, health care systems, and water systems."[19] While those quotes show modest room for scoping debate, there is clearly a well-developed common sense of what systems we simply cannot do without. Unfortunately, the quotes also suggest a common misunderstanding of what CI actually means and what systems might be affected when CI-specific regulations are passed.

Statutorily, the most common definition of "critical infrastructure" is found in the USA PATRIOT (Uniting and Strengthening America by Providing Appropriated Tools Required to Intercept and Obstruct Terrorism) Act of 2001:

> systems and assets, whether physical or virtual, 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.[20]

Notable here is the definition's immense flexibility. Because the definition fails to specify or limit what is "critical," the meaning of the term can be extended to include almost anything. Naturally, since 2001, that is exactly what has happened. As the Congressional Research Service reports, unconstrained by statute, CI has since strayed from an "earlier emphasis on the physical foundations of national power, to a wider concern with provision of essential services and customary conveniences to the public." [21]

What the Congressional Research Service's analysis speaks to is the breadth of the current system formally laid out in a presidential policy directive in 2013.[22] Today, critical infrastructure is divided between 16 official sectors, ranging from the clearly essential, including energy and water, to the decidedly optional, such as commercial facilities. Within each, we find boggling scale. In its profile of the chemicals sector's critical infrastructure, for instance, the Cybersecurity and Infrastructure Security Agency (CISA) lists a wide range of industries that by no definition should be considered critical, including cosmetics, perfumes, bookbindings, and vehicle paint.[23] The risk of a scratched car, it seems, is a matter of national security.

Across other sectors, we find similar breadth. Under transportation, there are critical systems such as our trains, but also vanpool and rideshare services.[24] Under commercial facilities, we find perhaps the biggest sprawl of "optionals," including the nation's 2.1 million office buildings and retail shopping centers as well as the entire hotel, film, broadcast, and casino industries.[25]

To paint this policy bloat in even sharper relief, not only are these sectors and their components sprawling in scope, but also they are sprawling in economic size. According to CISA documents, the combined GDP of just 3 of these 16 sectors—the chemicals sector (25 percent), commercial facilities sector (20 percent), and healthcare sector (17.4 percent)—represent over 50 percent of the total US economy.[26] While similar sector-specific figures are missing in the documents of many of the remaining 13 sectors, given the sweep of these categories and the major economic categories the other sectors represent, it is easy to imagine CISA's critical infrastructure designation covers a supermajority of US economic activity.

The best answer to "What is critical infrastructure?" is, it seems, another question: "What *isn't* critical infrastructure?"

To avoid potentially undue critique, we should note why this categorical sprawl has evolved. Fundamentally, our current CI policy frame was built to service not regulation but rather organized information sharing. If we look to the text of Presidential Policy Directive 21, the explicit

FIGURE 1. Critical infrastructure risk management framework

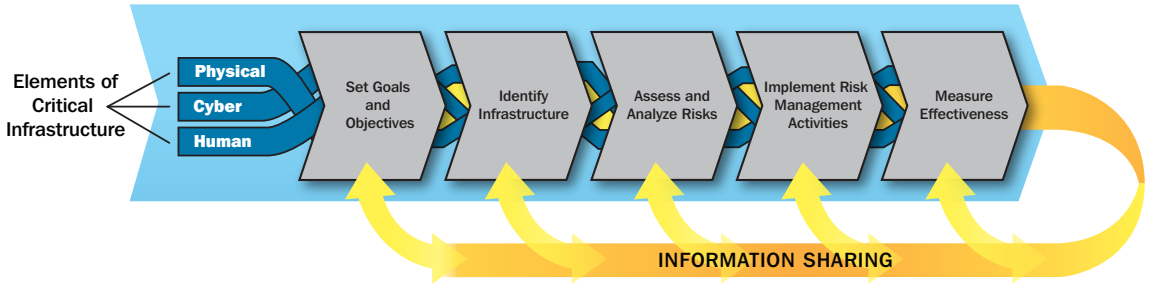

Note: Figure shows the goals of critical infrastructure policies as laid out in Presidential Policy Directive 21.

Source: Department of Homeland Security, *NIPP 2013: Partnering for Critical Infrastructure Security and Resilience* (Washington, DC: Department of Homeland Security, 2013), 15.

goals of our national CI policy stress enabling "efficient information exchange" and facilitating information "integration and analysis" to inform critical infrastructure decision-making (see also figure 1).[27] Given that these sectors were built to service threat analysis, information aggregation, and intelligence sharing, this overinclusion starts to carry a certain logic. With a broad net, agencies perhaps can create lines of threat communication across a greater swath of industries while also broadening the diversity of data sets and easing information gaps. In many ways, that design reflects the PATRIOT Act origins of critical infrastructure policy: A core problem in 2001 was a failure of imagination and a failure to connect the dots.[28] As a result, CI policy reflects a drive to de-silo information en masse and spot threats before it's too late.

The emerging challenge today, however, is that legislators have begun grafting regulation onto this structure originally designed for post-9/11 information sharing. In 2022, as a reaction to the Colonial Pipeline disaster, Congress passed the Cyber Incident Reporting for Critical Infrastructure Act (CIRCIA) of 2022. Explicitly regulatory, the act grants CISA new cybersecurity reporting regulatory power over all entities "in a critical infrastructure sector, as defined in Presidential Policy Directive 21"—that is, the businesses and organizations within those 16 categories.[29] While many likely saw the bill as narrowly targeted, the scope of critical infrastructure means it amounts to a broad, nearly economy-wide grant of cyber-reporting authority.

Thankfully, in the specific case of cyber reporting, the unwieldly breadth of this nebulous structure, while flawed, may be low impact. Rules to narrow the scope are under way, and given the nature of CIRCIA and cyber-reporting policy, risks are minimal. If this framework is pushed to accommodate more impactful regulations to rein in AI or any other potential risk, however, Congress and the executive could easily breach the load-bearing capacity of this ill-defined construction.

3. A Shaky Foundation for Regulation

For AI regulation, or any other expansion of CI regulations for that matter, what challenges does the breadth of the current policy frame present?

The first challenge is prioritization. If 50 percent or more of the economy falls under any new CI regulation, perhaps increasingly limited budgetary capacity naturally demands administrators set priorities.[30] What those priorities are, however, is unbounded and unclear. In October's AI executive order, we already see an early example of this prioritization question.[31] In the executive order, the Department of Homeland Security (DHS) is newly mandated to translate the National Institute of Standards and Technology's AI Risk Management Framework into "relevant safety and security guidelines for use by critical infrastructure owners and operators." Given the scope of CI's 16 sectors, this presidential ask lacks specificity. The DHS could go nearly any direction, and while some might assume the agency would naturally focus on the clearly critical sectors, such as the grid, current CI programming suggests it's just as likely that implementation will reflect the specific hobbyhorse priorities of those who happen to be in charge. Scanning current CI projects, we already find a varied picture, including initiatives on autonomous vehicle security, crowd control and safety, and buoy efficacy assessments.[32] While those topics may indeed be worthy of attention, their systemic criticality is questionable.

In the case of this benign DHS requirement, the prioritization problem may simply result in a mismatch between White House expectations and implementation reality. In the case of regulation, however, such government by hobbyhorse risks sending congressional legislation off course. Our grid and water infrastructure are indeed insecure; however, we risk failing to resolve these issues if the regulatory foundation doesn't set priorities.

That leads to a second risk: clarity. Because the bucket of officially critical infrastructure is so large and prioritization decisions must naturally be made, any industry caught in this regulatory snare faces immense uncertainty about whether it will face regulation. This is already a clear issue. Entering 2024, CISA is actively writing the implementing rules for the critical infrastructure cyber-reporting powers granted to it by CIRCIA. Likely, the rules will be tailored somewhat toward select targets. What those targets might be, however, is unclear. In public comments submitted in response to a CISA request for information, many industry commenters have highlighted the challenging uncertainty of the act's scope, with several urging clarification about who will be subject to reporting and when.[33]

While the uncertainty of the scope of reporting requirements is currently uncomfortable, if any future AI regulatory bills go further—perhaps mandating certain standards, designs, security controls, or behaviors—such uncertainty can be disabling. The website CIO reports that even without a bill, existing AI regulatory uncertainty has led 44 percent of large companies to take a "short pause" on AI deployment decisions.[34] If an unconstrained CI regulatory bill were ever passed, such numbers would likely skyrocket as industry waits to hear who might fall under the regula-

tory hammer. How long might that uncertainty last? In the case of CIRCIA, rules have yet to be made even two years after the bill's passage. If Congress passes any bill to regulate the use of AI in CI, it invites similar years-long pauses, during which investment will slow, diffusion will cease, and American competitiveness will be harmed.

The final (and perhaps greatest) challenge: overregulation. While it's somewhat safe to assume prioritization is necessary, the scope of CISA's CI designation opens the door to the opposite, disorder, as any bill passed to regulate the use of AI in CI is potentially tantamount to an economy-wide catchall AI regulation. Turning back to the recently introduced Thune-Klobuchar AI bill, we see an illustration of the risk of such legislation. The intent of the authors of the bill appears to be restraint; most, including industry, see it as a light-touch, targeted proposal, and that supposed conservativism has sparked its relative momentum. In the bill's text, the Department of Commerce is given new reporting and standards enforcement powers, narrowly targeted at the AI systems used in critical infrastructure that have a "legal or similarly significant effect on the direct management and operation of critical infrastructure" (borrowing the traditional PATRIOT Act definition). Supposedly, nearly all other systems will be free from regulation, and innovation can continue.

Even with the added "direct management and operation" qualifier, however, the Thune-Klobuchar bill's scope appears to cover many of the most compelling AI use cases today. Ride-sharing services, for instance, are counted by the Department of Transportation as part of the transportation CI sector. No doubt a Lyft driverless vehicle would qualify as "directly operated" by AI and therefore be subject to this regulation. Data centers, which currently slot into the information technology critical infrastructure sector,[35] likewise are often operated by AI systems geared at managing server loads, cooling, and other key operational services. Beyond those examples, we can imagine many other increasingly AI-operated services falling under CI, including software-defined networks, factory automation, hospitals and medical equipment, mining rigs, city buses, trains, drones, and likely many other applications. Today, AI is in active use in all those sectors and, in many cases, appears to fall within the regulatory bounds of this bill. Even if the intent of the authors is narrowness, because of the creaky regulatory foundation the text is built on, the authors open the regulatory door to so much more.

While this is just one bill, any other legislative concept that uses its foundation will face this same challenge. Critical infrastructure as commonly conceived may indeed be a reasonable regulatory target. Critical infrastructure as currently defined is not. If we want AI safety, secure systems, and continued innovation, modest work will be needed on the part of Congress.

4. Doing Better

Thankfully for AI policy, this problem has already been partially acknowledged by a small collection of cybersecurity policy experts. In its 2020 final report, the Cyberspace Solarium Commission,

a congressionally mandated body, proposed a more focused critical infrastructure categorization: systemically important critical infrastructure.[36] The intent here is a far more discrete list of only the most critical systems and entities. To avoid re-creating the bloat of the current system, the commission proposed a list of requirements, recently put in more actionable formulaic terms by the RAND Corporation, to narrow eligible assets and truly identify the systems we can't live without.[37]

While there is certainly room for debate on the scope of those example models, such narrowed precision would be a regulatory breath of fresh air. For any regulation, RAND's formulas would ensure targeted prioritization of assets and the elimination of the overregulation problem, while ensuring clarity about who and what might be subject to current and future CI regulation.

It's important to stress that action must come from Congress for these ideas to be implemented. Recently, CISA has expressed modest interest in using executive discretion to create a similarly narrowed list.[38] That is certainly a welcome step, yet a solution that is far too administratively flexible to contain mission creep over time and stop overregulation risks. Only through congressional action can we set firm legal boundaries, and only through Congress can we ensure that the regulations that already exist, such as CIRCIA, are built on this new structure. Because CI regulation is so new and any AI regulation has yet to pass, such changes may indeed be legislatively possible and should be included in any potential AI bill that targets critical infrastructure.

As we start 2024, we are riding on a clear wave of nonstop AI innovation. This technology has amazing potential to transform the United States and unleash abundance. Such promise demands, therefore, that we get any AI regulatory efforts right. Through modest changes, the massive scope of critical infrastructure can be focused, creating a foundation narrowly targeted at just our most critical systems. The result will be better administration, better safety, and the light touch that many in Congress are seeking. For the rest of the economy, AI can be unleashed, free to be diffused, used, and hopefully transformative.

ABOUT THE AUTHOR

Matthew Mittelsteadt is a research fellow and technologist for the AI and Progress Project whose work highlights the importance of AI diffusion and of ensuring that emerging AI technologies yield a net benefit. His research focuses on AI regulatory design and measurement, critical infrastructure and cybersecurity, and the unleashing of rapid AI diffusion. As a scholar, he places special emphasis on grounding policy making through education and helping policymakers understand AI through seminars and Mercatus's *AI Policy Guide*.

Mittelsteadt's work has been published by *The Hill*, *Noema Magazine*, and the Federal Judicial Center. It has been cited by media including the *New York Times*, *Bloomberg*, *Foreign Policy*, and *Politico*. Before joining Mercatus, he was a fellow at the Institute for Security Policy and Law,

where he focused on the intersection of AI and national security. He holds a BA in economics from St. Olaf College, an MPA from Syracuse University, and an MS in cybersecurity from New York University.

NOTES

1. Gabby Miller, "Transcript: House Hearing on DHS and CISA Role in Securing AI," *Tech Policy Press*, December 15, 2023.
2. Gabby Miller, "US Senate AI 'Insight Forum' Tracker," *Tech Policy Press*, December 8, 2023.
3. Rebecca Klar, "Thune, Klobuchar Release Bipartisan AI Bill," *The Hill*, November 15, 2023.
4. The White House, "Fact Sheet: Biden-Harris Administration Announces Key AI Actions Following President Biden's Landmark Executive Order," press release, January 29, 2024; Department of Homeland Security, "DHS/CISA and UK NCSC Release Joint Guidelines for Secure AI System Development," press release, November 26, 2023.
5. Dan Goodin, "Disaster Averted: Computer Intruder Tried to Poison Florida City's Drinking Water with Lye," *Ars Technica*, February 8, 2021.
6. Andy Greenberg, "The Colonial Pipeline Hack Is a New Extreme for Ransomware," *Wired*, May 8, 2021.
7. Sarah Braithwaite, "Denmark Faces Largest Cybersecurity Incident to Date," *Cybersecurity* (blog), University of Hawai'i–West O'ahu, December 8, 2023.
8. Jamie Tarabay and Katrina Manson, "Iranian Linked Hacks Expose Failure to Safeguard US Water System," *Bloomberg*, December 22, 2023.
9. Michael Martina et al., "US Officials Deliver Warning That Chinese Hackers Are Targeting Infrastructure, *Reuters*, January 31, 2024.
10. Lewis Milford and Shelley Robbins, "Texas Power Outage Deaths: Is Cruelty and Neglect Our New Energy Policy?," *The Hill*, June 28, 2021.
11. Mose Buchele, "One Year Later, Many Question the 'Official' Number of Deaths Linked to the Texas Blackout," *KUT News*, February 15, 2022.
12. Vida Rozite et al., "Why AI and Energy Are the New Power Couple," *IEA 50* (blog), International Energy Agency, November 2, 2023.
13. Sims Witherspoon and Will Fadrhonc, "Machine Learning Can Boost the Value of Wind Energy," *The Keyword* (blog), Google, February 26, 2019.
14. Alvaro Velasquez, "Guaranteeing AI Robustness against Deception (GARD)," Defense Advanced Research Projects Agency, accessed April 1, 2024, https://www.darpa.mil/program/guaranteeing-ai-robustness-against-deception.
15. Apostol Vassilev et al., *Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations*, NIST AI 100-2e2023 (Washington, DC: National Institute of Standards and Technology, January 2024).
16. Micah Musser, "Adversarial Machine Learning and Cybersecurity: Risks, Challenges, and Legal Implications," Center for Security and Emerging Technology, April 2023.
17. *Considering DHS' and CISA's Role in Securing Artificial Intelligence: Hearing before the Subcomm. on Cybersecurity and Infrastructure Protection of the H. Comm. on Homeland Security,* 118th Cong. (2023).
18. US Government Accountability Office, *Critical Infrastructure Protection: CISA Should Improve Priority Setting, Stakeholder Involvement, and Threat Information Sharing*, GAO-22-104279 (Washington, DC, March 2022), 1.

19. The White House, "A Proclamation on Critical Infrastructure Security and Resilience Month, 2023," press release, October 31, 2023.

20. Uniting and Strengthening America by Providing Appropriate Tools Required to Intercept and Obstruct Terrorism (USA PATRIOT) Act, Pub. L. No. 107-56, § 1016(e) (2001).

21. Brian E. Humphreys, *Critical Infrastructure: Emerging Trends and Policy Considerations for Congress*, CRS Report R45809 (Washington, DC: Congressional Research Service, July 8, 2019), 2.

22. The White House, "Presidential Policy Directive: Critical Infrastructure Security and Resilience," press release, February 12, 2023.

23. Cybersecurity and Infrastructure Security Agency, "Chemical Sector Profile," March 23, 2022.

24. Cybersecurity and Infrastructure Security Agency, "Transportation Systems Sector," accessed March 25, 2024, https://www.cisa.gov/topics/critical-infrastructure-security-and-resilience/critical-infrastructure-sectors/transportation-systems-sector.

25. Cybersecurity and Infrastructure Security Agency, "Commercial Facilities Sector," accessed March 25, 2024, https://www.cisa.gov/topics/critical-infrastructure-security-and-resilience/critical-infrastructure-sectors/commercial-facilities-sector.

26. Cybersecurity and Infrastructure Security Agency, "Chemical Sector," accessed March 25, 2024, https://www.cisa.gov/topics/critical-infrastructure-security-and-resilience/critical-infrastructure-sectors/chemical-sector; Cybersecurity and Infrastructure Security Agency, "Commercial Facilities Sector"; US Department of Homeland Security, *Healthcare and Public Health Sector-Specific Plan* (Washington, DC, May 2016), 4.

27. The White House, "Presidential Policy Directive."

28. Thomas H. Kean et al., *The 9/11 Commission Report: Final Report of the National Commission on Terrorist Attacks upon the United States* (Washington, DC: Government Printing Office, 2004).

29. Cyber Incident Reporting for Critical Infrastructure Act (CIRCIA) of 2022, Pub. L. No. 117-103, div. Y. The full text of the act is available at https://www.congress.gov/bill/117th-congress/house-bill/2471/text.

30. Christian Vasquez, "CISA Budget Cuts Would Be 'Catastrophic,' Official Says," *Cyberscoop*, October 25, 2003.

31. Justin Hendrix, "White House Executive Order on AI Gives Sweeping Mandate to DHS," *Tech Policy Press*, November 1, 2023.

32. Ashley Albrecht, "CIRI Connects with Automakers at Auto-ISAC Cybersecurity Summit," October 8, 2018, Critical Infrastructure Resilience Institute, University of Illinois, Urbana-Champaign; Dimitri Kusnezov, "Protecting Our Critical Infrastructure during Uncertain Times," Science and Technology Directorate, Department of Homeland Security, press release, November 1, 2023.

33. See, for example, the comments at https://www.regulations.gov/docket/CISA-2022-0010/comments?filter=Comment%20CISA-2022-0010%20.

34. Maria Korolov, "Regulatory Uncertainty Overshadows Gen AI Despite Pace of Adoption," *CIO*, August 24, 2023.

35. Cybersecurity and Infrastructure Security Agency, "Information Technology Sector," accessed March 25, 2024, https://www.cisa.gov/topics/critical-infrastructure-security-and-resilience/critical-infrastructure-sectors/information-technology-sector.

36. Cyberspace Solarium Commission, *Final Report of the Cyberspace Solarium Commission* (Washington, DC: Government Printing Office, March 2020).

37. Cyberspace Solarium Commission, *Legislative Proposals* (Washington, DC: Government Printing Office, July 2020); John Bordeaux et al., *Identifying and Prioritizing Systemically Important Entities: Advancing Critical Infrastructure Security and Resilience* (Homeland Security Operational Analysis Center, RAND, 2023).

38. Justin Doubleday, "CISA Establishing 'Systemically Important Entities' Office," *Federal News Network*, March 23, 2023.