How Might China Conduct AI-Enabled Cyber Warfare Against the Homeland?

KEY TAKEAWAYS

- CSIS researchers ran a war game on AI-enabled cyber threats to critical infrastructure for members of the House Committee on Homeland Security and the House Select Committee on China. The bipartisan group of members played the United States, and an AI agent created specifically for the game played the People's Republic of China (PRC).
- In the scenario, Beijing plans to attack Taiwan. It uses AI-enabled cyberattacks to disrupt critical infrastructure inside the United States, hoping to dissuade or prevent the United States from coming to Taiwan's aid.
- The AI agent CSIS created to play Beijing attacked first, going after port infrastructure. As the members playing the United States responded, the AI agent playing Beijing learned and adapted, shifting to go after undefended infrastructure and eventually to poison U.S. AI systems.
- Members examined how the targeting of freight rail, port terminals, transportation networks, and utilities, along with AI models trained on vulnerable data sets, could challenge the United States' defensive posture and result in cascading impacts on the preparedness and resilience of the homeland.

By Benjamin Jensen and Yasir Atalan of the Futures Lab and Emily Harding of the Intelligence, National Security, and Technology Program

BACKGROUND AND CONTEXT

Previous work by CSIS's **Futures Lab** and **Intelligence, National Security, and Technology Program** (INT) has highlighted both the vulnerability of U.S. critical infrastructure and Beijing's willingness to attack these civilian targets to accomplish military goals. In ***A Playbook for Winning the Cyber War***, INT outlined how China **operates** in the cyber domain and how the United States is unprepared to defend itself. In an **earlier iteration** of this war game, researchers demonstrated that policymakers need to think differently about the future of cyber warfare if the United States is to match and defeat adversaries like China.

The new war game used AI in two ways. First, CSIS created a custom AI agent named Red Kraken to play China. The Futures Lab trained the agent on Chinese doctrine and a dataset of over 150 previously documented PRC cyber operations, tested the agent over 100 times, and deployed it to play "red" in the game. The Kraken attacked and reacted to U.S. moves. As a trained agent, Kraken can avoid the bias of a U.S. player attempting to think like the adversary.

Second, CSIS assumed that China and the United States both had Mythos-level capabilities in AI, allowing each side to attack the other using AI agents. Adopting these tools also means defending them, along with the data that underpins them, against infiltration and poisoning of that data. This capability inclusion led to an increased attack surface in addition to stronger offensive capabilities for each side.

China's theory of victory in both games involved distracting the United States with chaos at home, which would prevent or at least delay a U.S. deployment to the Pacific to aid Taiwan. During the delay, while the United States is attempting to bring its port infrastructure back online and untangle its logistics networks, China could establish a foothold in Taiwan.

LEGISLATIVE AND POLICY IMPLICATIONS

Several bills that seek to strengthen cybersecurity for the homeland are in progress. The **Guaranteeing Universal Access to Cybersecurity Act** (S. 4699), introduced in early June 2026, would fund the Multi-State Information Sharing and Analysis Center (MS-ISAC). The MS-ISAC allows states to share cybersecurity threat information among themselves. The **PILLAR Act** (H.R. 5078) would extend

RECOMMENDATIONS

- Congress should provide authorities and funding to harden critical infrastructure, particularly at the state and local levels. Matching grants for state efforts can encourage state and local governments to grow their own capabilities.
- Apply AI agents to act as the "red" side in a variety of national security contexts. If properly trained, AI can help policymakers think like the adversary more accurately, removing personal and cultural assumptions. The intelligence community also should experiment with using AI agents to play the "red" side in order to test their own assumptions about how an adversary might act.
- Evaluate steps to safeguard the data that underpins AI tools. These are potential vulnerabilities, exploited by Red Kraken in this exercise, that could undermine entire logistics chains in a critical time.

the State and Local Cybersecurity Grant Program through FY 2035 and modernize its scope, helping to fund cybersecurity below the federal level.

The **Cybersecurity Information Sharing Act of 2015** gave companies liability protection and antitrust safeguards for voluntarily sharing cyber threat indicators with the government and each other. The law sunset on September 30, 2025, and has survived on a string of short-term patches. The provisions were most recently extended to September 30, 2026, as part of the FY 2026 Consolidated Appropriations Act.

CHALLENGES AND RISKS

- Defense of critical infrastructure in the homeland is absolutely essential, and it must be done before a conflict begins. China and other sophisticated cyber adversaries understand that lightly defended U.S. infrastructure is a vulnerability.
- As China and other adversaries develop AI agents like Mythos, they will be able to target these vulnerabilities faster and at a larger scale.
- Simultaneously, as the United States seeks to lead the world in the development and adoption of AI, labs need to have a security mindset. The algorithms must be thoroughly tested and the data sets on which they are trained must be carefully secured.

ADDITIONAL RESOURCES

Emily Harding, Julia Dickson, and Aosheng Pusztaszeri, *A Playbook for Winning the Cyber War* (Washington, DC: CSIS, September 2025), **https://www.csis.org/programs/intelligence-national-security-and-technology-program/playbook-winning-cyber-war**.

Representative Michael McCaul, "Opening Remarks: Congressional Wargame on AI-Enabled Cyber Threats to U.S. Critical Infrastructure," CSIS, July 21, 2026, **https://www.csis.org/events/opening-remarks-congressional-wargame-ai-enabled-cyber-threats-us-critical-infrastructure**.

Yasir Atalan, "What to Know About Chinese AI Models," CSIS, *Critical Questions*, July 2, 2026, **https://www.csis.org/analysis/what-know-about-chinese-ai-models**.

Benjamin Jensen and Yasir Atalan, "The Escalation Danger Trump and Xi Must Address," CSIS, *Commentary*, May 12, 2026, **https://www.csis.org/analysis/ai-escalation-danger-trump-and-xi-must-address**.

Benjamin Jensen, Yasir Atalan, and Dan Tadross, "It Is Time to Democratize Wargaming Using Generative AI," CSIS, *Commentary*, February 22, 2024, **https://www.csis.org/analysis/it-time-democratize-wargaming-using-generative-ai**.

Yasir Atalan, Benjamin Jensen, and Jose M. Macias III, *Eroding Trust in Government: What Games, Surveys, and Scenarios Reveal about Alternative Cyber Futures* (Washington, DC: CSIS, April 2024), **https://www.csis.org/analysis/eroding-trust-government-what-games-surveys-and-scenarios-reveal-about-alternative-cyber**.

Benjamin Jensen, Devi Nair, Yasir Atalan, and Jose M. Macias III, *CISA's Evolving.gov Mission: Defending the United States' Federal Executive Agency Networks* (Washington, DC: CSIS, October 2023), **https://www.csis.org/analysis/cisas-evolving-gov-mission-defending-united-states-federal-executive-agency-networks**.

Brandon Valeriano, Benjamin Jensen, and Ryan Maness, *Cyber Strategy: The Evolving Character of Power and Coercion* (New York: Oxford University Press, 2018), **https://global.oup.com/academic/product/cyber-strategy-9780190618094?cc=us&lang=en&**.

For more information, contact **Chloe Himmel** at **202.775.3186** or **chimmel@csis.org**.