CENTER FOR CYBERSECURITY POLICY AND LAW

WHITEPAPER

Playbook for Expanding Cyber Resilience in the Age of Frontier AI Models

SEPTEMBER 2026

Table of Contents

Executive Summary

Frontier artificial intelligence (AI) is transforming cybersecurity. Advanced AI models from OpenAI and Anthropic as well as other similar frontier systems are changing how cyber threats are developed, how defenders respond, and how governments must think about national resilience. Open-weight models are also rapidly advancing, now trailing the most advanced closed-weight models on cyber capability evaluations by just roughly five months. These technologies enable malicious actors to automate cybercrime, improve social engineering, accelerate vulnerability discovery, and conduct increasingly sophisticated influence operations. At the same time, they are providing defenders with new capabilities to strengthen threat intelligence, improve incident response, augment limited cybersecurity workforces, and enhance decision-making across government.

For governments operating with constrained cybersecurity resources, these developments present both significant challenges and unprecedented opportunities. Limited budgets, workforce shortages, dependence on foreign technology providers, and competing national priorities often make it difficult to build and sustain mature cybersecurity capabilities. Frontier AI has the potential to help address some of these constraints by increasing the effectiveness of existing institutions and enabling governments to accomplish more with limited resources. However, AI is not a substitute for sound cybersecurity governance, skilled professionals, or sustained investment.

This paper argues that the emergence of frontier AI does not fundamentally change the principles of an effective national cybersecurity strategy (NCS); rather, it changes the environment in which those strategies are developed and implemented. Governments should continue to build cybersecurity strategies around risk assessment, clear strategic objectives, governance, workforce development, critical infrastructure protection, public-private collaboration, and continuous evaluation. Certain foundational priorities, such as modernizing vulnerability management and reducing legacy technology debt, warrant immediate attention regardless of a nation's AI adoption timeline.[1] Frontier AI should be considered alongside these priorities not as an objective, but as a tool that may strengthen their implementation where it provides measurable value.

Importantly, this paper does not suggest that every nation should make significant investments in frontier AI or seek to develop sovereign frontier AI capabilities. Countries differ substantially in their levels of digital maturity, available financial resources, technical expertise, and national priorities. For some governments, investments in foundational cybersecurity capabilities may provide greater near-term benefits than investments in advanced AI technologies. For others, frontier AI may serve as a force multiplier that accelerates cybersecurity capacity and improves operational effectiveness. The appropriate path will differ from country to country.

[1] https://www.nsa.gov/Press-Room/News-Highlights/Article/Article/4523810/five-eyes-cyber-security-agencies-statement/

Ultimately, the key policy question is not whether governments should adopt frontier AI, but how frontier AI changes the decisions they must make. Every nation should assess how AI is reshaping its cyber risk environment, determine where AI can most effectively support national cybersecurity objectives, and continually evaluate whether the benefits of adoption outweigh the associated costs, dependencies, and risks.

Countries that approach frontier AI through the same disciplined process used to develop an effective national cybersecurity strategy, one grounded in governance, partnerships, risk management, and long-term resilience, will be best positioned to protect their citizens, strengthen their digital economies, and adapt to an increasingly AI-enabled cyber landscape.

The Center has been engaging with many stakeholders and has compiled the lessons learned into this playbook. We want to thank Cisco for funding this paper, as well as contributing their knowledge and expertise to the final product.

1. The AI Revolution and National Cybersecurity

Artificial intelligence is reshaping nearly every aspect of the global economy, national security, and digital society. While previous generations of AI largely automated narrow, well-defined tasks, today's frontier AI models are capable of reasoning across complex problems, generating sophisticated software, analyzing vast quantities of information, interacting with external tools, and increasingly acting as autonomous agents.[2] [3] These capabilities are rapidly transforming how governments, businesses, and individuals use technology. They are also fundamentally changing the cybersecurity landscape.[4]

Frontier AI refers to the most advanced generation of general-purpose artificial intelligence models developed at the leading edge of research.[5] These models are trained on enormous datasets and possess broad capabilities that extend well beyond text generation.[6] They can write and analyze software, discover patterns across massive datasets, summarize intelligence, reason through technical problems, assist with scientific research, and increasingly execute complex workflows with minimal human supervision. As these models continue to improve, they are becoming foundational technologies that will influence virtually every sector of society, much as the internet and cloud computing did in previous decades.

---

[2] https://www.frontiermodelforum.org/technical-reports/managing-advanced-cyber-risks-in-frontier-ai-frameworks/

[3] https://aws.amazon.com/blogs/aws-insights/the-rise-of-autonomous-agents-what-enterprise-leaders-need-to-know-about-the-next-wave-of-ai/

[4] https://www.anthropic.com/glasswing

[5] https://www.crowdstrike.com/en-us/cybersecurity-101/artificial-intelligence/frontier-ai/

[6] https://www.anthropic.com/claude/mythos

The emergence of frontier AI represents a strategic inflection point because it simultaneously lowers the barriers to sophisticated cyber operations while dramatically expanding the capabilities available to defenders. Historically, advanced cyber capabilities required highly specialized expertise, significant computing resources, and years of operational experience. Today, many of those capabilities are becoming increasingly accessible through AI-assisted tools that can help identify software vulnerabilities, generate code, automate analysis, accelerate research, and support operational decision-making.[7] Although human expertise remains essential, frontier AI is compressing the time, cost, and expertise traditionally required to perform many cybersecurity tasks.

The emergence of autonomous, agentic AI-orchestrated cyber operations illustrates how quickly this threat is evolving.[8] Malicious actors can leverage frontier AI to automate phishing campaigns, generate convincing social engineering content in dozens of languages, develop malware more rapidly, identify exploitable vulnerabilities, conduct reconnaissance, and produce realistic synthetic media capable of undermining public trust. As these capabilities become more widely available, cybercrime, espionage, and influence operations are likely to become faster, cheaper, and more scalable than ever before.[9] Rather than creating entirely new categories of cyber threats, frontier AI accelerates existing threats by increasing their speed, sophistication, and accessibility, as demonstrated by recent state-linked, AI-orchestrated cyber campaigns.

At the same time, the very capabilities that benefit attackers can provide equally significant advantages for defenders. AI is increasingly enabling cybersecurity teams to detect threats earlier, analyze malware more efficiently, prioritize vulnerabilities based on risk, automate routine security operations, and accelerate incident response. National cybersecurity centers can use AI to process vast quantities of threat intelligence that would otherwise overwhelm human analysts. Governments can employ AI to improve risk assessments, identify emerging trends, model potential attack scenarios, and better allocate scarce cybersecurity resources.[10] [11] Frontier AI therefore has the potential to become an important force multiplier that allows even relatively small cybersecurity organizations to operate with greater speed, consistency, and effectiveness.

The increasing integration of AI into critical infrastructure, government services, defense systems, healthcare, financial services, telecommunications, and industrial operations also introduces an entirely new set of security considerations.[12] AI systems themselves become high-value targets for

[7] https://www.interpol.int/en/Resources/INTERPOL-Spotlight/Spotlight-Issue-2-Cybercrime/Spotlight-Cybercrime-Innovation

[8] https://www.anthropic.com/news/disrupting-AI-espionage

[9] https://www.europol.europa.eu/publications-events/publications/malicious-uses-and-abuses-of-artificial-intelligence

[10] https://www.cyber.gov.au/sites/default/files/2026-05/Opportunities%20for%20AI%20in%20cyber%20defence.pdf

[11] https://www.cisa.gov/sites/default/files/2025-01/JCDC%20AI%20Playbook_FACT%20SHEET.pdf

[12] https://www.worldbank.org/en/results/2025/01/29/enhancing-cyber-resilience-in-developing-countries

adversaries seeking to manipulate training data, exploit model vulnerabilities, steal proprietary models, or influence AI-generated outputs. Protecting AI systems is therefore becoming inseparable from protecting national digital infrastructure.[13] Governments must increasingly consider AI security, governance, transparency, and resilience as core components of national cybersecurity rather than as separate technology policy issues.

These developments carry particular significance for governments with constrained cybersecurity resources. Many countries continue to face persistent shortages of cybersecurity professionals, limited government budgets, fragmented institutional responsibilities, and growing dependence on foreign technology providers. At the same time, digital transformation is accelerating across every region of the world. Governments are expanding digital public services, modernizing critical infrastructure, digitizing financial systems, and connecting more citizens and businesses to the internet.[14] This growing digital dependence increases both economic opportunity and national cyber risk.[15] [16]

This shift is not limited to frontier models alone. A broader AI cybersecurity ecosystem is emerging, including smaller open-weight models, specialized security harnesses, and coding-assistance tools that complement rather than compete with frontier systems.[17] [18] [19] These alternatives can be particularly valuable for governments with constrained budgets or sovereignty concerns about granting external providers access to sensitive government environments.

For the purposes of this paper, "smaller nations" refers broadly to governments operating with constrained cybersecurity resources, whether due to budget, workforce, institutional capacity, digital maturity, or access to advanced technologies. It is not intended as a measure of geography, population, economic size, or national importance. Throughout this paper, the term is used to describe governments that must make difficult strategic tradeoffs when allocating limited cybersecurity resources.

Traditionally, building mature national cybersecurity capabilities required decades of investment in specialized institutions, highly trained personnel, and sophisticated technical infrastructure. Frontier AI has the potential to alter that equation. Properly governed and strategically deployed, AI can augment limited workforces, automate repetitive analytical tasks, improve policy development, enhance national cyber defense capabilities, and support more informed decision-

[13] https://www.cisa.gov/resources-tools/resources/principles-secure-integration-artificial-intelligence-operational-technology

[14] https://www.weforum.org/publications/digital-transformation-initiative/

[15] https://digiamericas.org/wp-content/uploads/2024/05/LATAM-CISO-REPORT-2024_.pdf

[16] https://digiamericas.org/wp-content/uploads/2026/05/SC_EN.pdf

[17] https://blogs.cisco.com/ai/introducing-antares-the-most-efficient-open-weight-ai-models-for-vulnerability-localization

[18] https://blogs.cisco.com/ai/announcing-foundry-security-spec

[19] https://github.com/project-codeguard

making across government. Rather than attempting to replicate the extensive cybersecurity institutions of larger powers, smaller nations have an opportunity to leapfrog traditional stages of cyber capacity development by integrating AI into both the development and implementation of their national cybersecurity strategies.[20]

This opportunity does not eliminate the need for strong governance, international partnerships, workforce development, or sustained investment. On the contrary, frontier AI increases the importance of these foundational elements. Nations that successfully combine sound cybersecurity governance with responsible AI adoption where necessary will be better positioned to strengthen national resilience, protect critical infrastructure, and accelerate digital development despite limited resources. Those that fail to adapt risk widening the gap between their digital ambitions and their ability to secure them.

2. How Frontier AI Changes the National Cybersecurity Landscape

AI Changes the Economics of Cybersecurity

Artificial intelligence is changing the economics of cybersecurity by altering the relationship between capability, cost, and scale. Historically, building sophisticated cybersecurity capabilities required significant financial investment, specialized technical expertise, and years of institutional development. Whether conducting advanced cyber operations or defending against them, capability generally grew in proportion to the number of skilled personnel and the resources available to support them.[21]

Frontier AI is beginning to change that relationship. Many tasks that previously required experienced analysts, including software analysis, vulnerability identification, malware investigation, policy research, and threat intelligence synthesis, can now be completed more quickly with AI-assisted tools. Likewise, malicious actors can use frontier AI to scale cyber operations at a meaningfully lower cost than traditional methods, though the most advanced frontier models still carry substantial operating costs. This dynamic may also create a cost asymmetry that favors attackers: adversaries can often achieve their objectives through discrete, wave-based attacks, whereas defenders must sustain continuous monitoring and response capabilities around the clock.

As a result, frontier AI is lowering the marginal cost of many cybersecurity activities for both attackers and defenders. Organizations can accomplish more with the same workforce, while

[20] https://www.centerforcybersecuritypolicy.org/insights-and-research/developing-a-national-cybersecurity-strategy
[21] https://www.sciencedirect.com/science/article/pii/S096386872400043X

adversaries can target more victims with fewer resources. The strategic advantage increasingly belongs not simply to those with the largest budgets, but to those that most effectively combine skilled personnel, sound governance, and AI-enabled capabilities.

For smaller nations, this shift presents both opportunity and risk. Frontier AI may allow governments with limited resources to strengthen certain cybersecurity functions without proportionally increasing staffing or budgets. At the same time, it reduces the barriers for adversaries to conduct sophisticated operations against a broader range of targets. Consequently, governments should evaluate frontier AI not as a standalone technology initiative, but as a factor that changes the return on investment of existing cybersecurity capabilities. Frontier AI should compete for cybersecurity investment alongside every other cybersecurity priority, not be exempt from the same cost-benefit analysis applied to any other national investment. In some cases, targeted investments in frontier AI may significantly improve national resilience. In others, investments in foundational capabilities such as governance, workforce development, asset management, or incident response may continue to deliver greater value. It is also important to distinguish between adopting frontier AI cybersecurity tools and pursuing industrial policy aimed at developing sovereign national AI models. The former is primarily a procurement and integration decision available to nearly any government today, while the latter is a far larger, longer-term undertaking with a different risk and cost profile. Nor should investment be limited to frontier AI alone: the broader AI cybersecurity ecosystem, including smaller models and purpose-built tools, often delivers comparable benefit at lower cost and complexity.

This changing economic landscape reinforces the idea that the objective of nations is not to adopt frontier AI for its own sake, but to determine where it provides the greatest strategic advantage in support of national cybersecurity objectives.

New Risks

Frontier AI does not fundamentally change the objectives of cyber adversaries: they continue to seek financial gain, espionage opportunities, disruption, or political influence. What changes is the speed, scale, and sophistication with which many of these activities can be conducted.

AI-Enabled Cybercrime

Cybercriminals have historically relied on automation to increase the scale of their operations. Frontier AI significantly expands those capabilities, and the emergence of autonomous, agentic AI operations shows that this expansion is no longer merely theoretical. AI-assisted coding tools can

also accelerate malware development, automate reconnaissance, and assist less experienced attackers in identifying exploitable vulnerabilities.[22]

This evolution lowers the barriers to entry for some forms of cybercrime while allowing sophisticated criminal organizations to operate more efficiently. Smaller nations may find themselves increasingly targeted not because they possess unique strategic value, but because automated attacks make it economically viable to target a much larger number of governments and organizations simultaneously.

Critical Infrastructure Risks

As AI becomes integrated into critical infrastructure like energy systems, telecommunications, transportation, healthcare, and financial services, it expands both operational capabilities and potential attack surfaces.[23] This expanded attack surface often results from new supply chain dependencies on AI model providers and from the external connectivity that AI systems require, which can force previously air-gapped or isolated operational environments to establish new network connections. AI may improve predictive maintenance, optimize industrial processes, or support real-time operational decisions. At the same time, AI-enabled systems become attractive targets for adversaries seeking to manipulate data, influence automated decision-making, or disrupt essential services.

Governments should recognize that adopting AI within critical infrastructure creates new cybersecurity responsibilities. Securing AI models, protecting training data, validating outputs, and maintaining human oversight become increasingly important components of critical infrastructure protection.

Supply Chain and Dependency Risks

Few smaller nations will develop frontier AI models domestically. Most will depend on commercial providers, cloud services, open-source models, or regional partnerships. While these approaches can significantly reduce costs and accelerate adoption, they also create strategic dependencies.

Governments should carefully evaluate where these dependencies introduce national risk. Questions surrounding data sovereignty, vendor concentration, cloud resilience, software supply chains, model transparency, and geopolitical considerations become increasingly important as AI systems are incorporated into government operations. For many countries, managing AI dependencies may prove more important than developing sovereign AI capabilities.

---

[22] https://www.interpol.int/en/Resources/INTERPOL-Spotlight/Spotlight-Issue-2-Cybercrime/Spotlight-Cybercrime-Innovation

[23] https://digiamericas.org/wp-content/uploads/2026/05/SC_EN.pdf

Frontier AI also changes the cost structure of cybersecurity. While some defensive capabilities become more accessible, access to the most advanced models may remain prohibitively expensive for many governments in the near term. Policymakers should therefore evaluate the total cost of ownership—including model access, compute, integration, governance, training, and ongoing operations—when determining whether frontier AI represents the best use of limited national cybersecurity resources.

New Opportunities

While much attention has focused on the risks associated with frontier AI, governments should devote equal attention to its potential benefits. For smaller nations in particular, AI may offer opportunities to augment limited human resources, improve decision-making, and increase operational efficiency without requiring proportional increases in staffing.

The objective should not be to automate cybersecurity indiscriminately, but rather to identify where AI provides measurable improvements to existing government capabilities.

Improved Cyber Defense

AI is increasingly capable of assisting cybersecurity professionals by analyzing logs, identifying anomalies, prioritizing vulnerabilities, and correlating threat intelligence across multiple data sources. Many of these activities require significant analyst time but involve repetitive or time-intensive work that AI systems can increasingly support.[24] [25] [26]

For governments facing persistent workforce shortages, AI may enable cybersecurity teams to focus their expertise on higher-value analytical and decision-making activities while reducing the burden of routine operational tasks.

Enhanced Threat Intelligence

National cybersecurity organizations routinely receive more information than analysts can reasonably process. AI can assist by organizing large volumes of threat reporting, identifying emerging trends, summarizing intelligence from multiple sources, translating foreign-language reporting, and highlighting information most relevant to national priorities. Rather than replacing intelligence analysts, AI enables them to devote more attention to strategic assessment, contextual analysis, and policy recommendations.

Greater Efficiency in Government Operations

[24] https://www.saif.google/secure-ai-framework

[25] https://www.cyber.gov.au/sites/default/files/2026-05/Opportunities%20for%20AI%20in%20cyber%20defence.pdf

[26] https://www.anthropic.com/research/mythos-preview

Beyond cybersecurity operations themselves, AI may improve many of the functions that support national cybersecurity. Governments can use AI to assist with drafting policies, analyzing legislation, summarizing stakeholder consultations, identifying regulatory gaps, evaluating international best practices, and supporting strategic planning.

These applications may be particularly valuable for smaller governments that have limited policy staff or rely on small teams responsible for multiple portfolios. In many cases, the most immediate return on AI investment may come not from advanced cyber operations but from improving the efficiency of government itself.

Selecting the Right AI Capability for the Task

One of the greatest misconceptions surrounding frontier AI is that the most advanced model is always the best solution for every cybersecurity problem. In practice, governments should match AI capabilities to operational requirements rather than assume that every task requires access to the most powerful, and often most expensive, models available.

Many cybersecurity activities can already be performed effectively using commercially available AI assistants, smaller language models, or AI-enabled cybersecurity products.[27] Tasks such as summarizing threat intelligence, drafting policies, analyzing legislation, translating technical documentation, generating cybersecurity awareness materials, reviewing security logs, or assisting with routine vulnerability management often benefit from AI capabilities that are widely accessible and comparatively inexpensive. These applications may provide immediate operational value while requiring relatively modest investments in infrastructure, training, and governance.

Frontier AI models offer additional capabilities for more specialized cybersecurity applications. They are increasingly demonstrating the ability to reason through highly complex technical problems, analyze sophisticated software, assist with advanced vulnerability research, and support dynamic analysis of compiled code. These capabilities may prove particularly valuable for organizations engaged in software development, advanced cyber defense, or vulnerability discovery. However, access to frontier models remains limited, their operational costs can be significant, and their advantages are often most pronounced in specialized use cases rather than routine cybersecurity operations.

Consequently, governments should avoid viewing frontier AI as the starting point for national cyber capability development. This does not mean governments should avoid frontier models altogether; rather, they should target frontier AI at a narrower set of high-value workloads while building broader capabilities using more accessible tools. Governments should first establish mature cybersecurity processes, develop institutional expertise in integrating AI into operational workflows, and gain experience using the AI technologies already available to them. As

[27] https://www.cisa.gov/ai/cisa-use-cases

organizational maturity increases and frontier AI becomes more accessible, governments will be better positioned to determine where these advanced capabilities provide sufficient operational benefit to justify their additional cost and complexity.

This graduated approach also reduces strategic risk. By developing model-agnostic workflows and integrating AI into existing cybersecurity processes, governments maintain flexibility as AI technologies continue to evolve. Frontier AI capabilities will almost certainly improve and become more broadly available over time, while smaller, less resource-intensive models will continue to advance as well. Organizations that build experience with AI today, regardless of the specific model they use, will be better prepared to evaluate and adopt future frontier capabilities as they mature.

Ultimately, the objective is not to deploy the most advanced AI model available. It is to improve national cybersecurity outcomes. Governments should therefore evaluate AI technologies using the same criteria they apply to any cybersecurity investment: whether they meaningfully reduce risk, strengthen resilience, improve operational effectiveness, and represent the best use of limited national resources.

Strategic Implications for Governments

The emergence of frontier AI requires governments to reconsider how they assess cyber risk, establish national priorities, and build long-term resilience. This does not necessarily mean that every nation should immediately invest in frontier AI capabilities. Rather, governments should evaluate how AI changes their operating environment and determine where targeted investments produce the greatest strategic value.

Changes to National Risk Assessments

National risk assessments should increasingly consider AI as both an enabling technology and a potential source of systemic risk. This includes evaluating AI-enabled cybercrime, synthetic media, dependencies on untrusted foreign AI providers, AI adoption within critical infrastructure, and the security of government AI deployments. Countries should also assess whether AI changes the likelihood or consequences of existing threats rather than treating AI as a separate category of risk.

Changes to Cybersecurity Priorities

Governments may find that AI shifts the relative importance of existing cybersecurity investments. For some countries, strengthening digital identity systems, improving information integrity, modernizing Security Operations Centers (SOCs), or expanding workforce training may provide greater value than investing directly in frontier AI technologies. Other priorities, including addressing legacy systems and technical debt, reducing attack surface, modernizing and

accelerating vulnerability management and patching, and securing open-source software libraries, warrant similar urgency to prepare against AI-enabled attacks.[28] [29]

Rather than asking, "How do we adopt AI?" policymakers should ask, "Where can AI most effectively advance our existing cybersecurity objectives?" In some cases, the answer may involve adopting commercial AI-enabled security tools. In others, it may involve investing first in foundational capabilities such as asset management, secure cloud adoption, or workforce development before pursuing more advanced AI initiatives.

Changes to Resilience Planning

National resilience planning must also evolve. Governments should prepare not only for AI-enabled cyber incidents but also for disruptions affecting AI systems themselves. Incident response plans, business continuity planning, crisis communications, and public trust mechanisms should account for new threats.

Importantly, resilience should not be measured by the extent to which a country adopts AI, but by its ability to continue delivering essential services regardless of how AI technologies evolve. For some nations, selective adoption supported by trusted commercial partners or regional cooperation may represent the most resilient and cost-effective path forward. For others with greater resources and more advanced digital ecosystems, deeper AI integration may be appropriate. The appropriate strategy will differ across countries, but every government should make that decision deliberately, based on national priorities, available resources, and an informed assessment of both risks and opportunities.

3. Building an AI-Era National Cybersecurity Strategy

Frontier AI does not eliminate the need for a national cybersecurity strategy; it reinforces its importance. As discussed in the Center for Cybersecurity Policy and Law's previous paper on developing a National Cybersecurity Strategy (NCS), every country should begin by understanding its own risks, establishing clear strategic objectives, designing effective governance, engaging stakeholders across society, and identifying the resources necessary to implement the strategy.[30] These principles remain as relevant in the age of AI as they were before the emergence of frontier models. What has changed is the environment in which those strategies must operate.

---

[28] https://fidoalliance.org/fido-alliance-digital-credentials/

[29] https://www.reuters.com/legal/litigation/ecb-tells-banks-draw-up-plans-against-ai-attacks-amid-disruption-fears-2026-07-07/

[30] https://www.centerforcybersecuritypolicy.org/insights-and-research/developing-a-national-cybersecurity-strategy

Governments should therefore resist the temptation to build separate AI-focused cybersecurity strategies. Instead, AI considerations should be integrated throughout an existing national cybersecurity strategy, ensuring that AI is evaluated alongside every other technological, economic, and national security development affecting the country's digital future. The appropriate role for AI will differ considerably across nations. Countries with advanced digital ecosystems and mature cybersecurity institutions may pursue extensive AI integration, while others may determine that investments in foundational cybersecurity capabilities should take precedence. Both approaches can be strategically sound if they are grounded in an objective assessment of national priorities, resources, and risk.

Adapting Existing Cybersecurity Strategies

An effective NCS begins with a comprehensive understanding of national risk. Governments should first evaluate how frontier AI changes their threat environment before determining whether AI investments are appropriate.

Rather than asking whether the country needs an AI strategy, policymakers should consider questions such as:

- How does AI change our national cyber risk profile?
- Which sectors are most likely to be affected by AI-enabled threats?
- Where could AI improve existing cybersecurity capabilities?
- What investments, including replacing legacy systems, improving resilience, and hardening existing infrastructure, rather than simply adopting new AI capabilities, would provide the greatest return given our available resources?
- What dependencies on foreign AI providers or cloud infrastructure introduce strategic risk?

The answers to these questions will vary significantly across countries. A nation with an established SOC, mature incident response capabilities, and a sizable cybersecurity workforce may identify operational AI as a high priority. Another nation may conclude that improving asset management, implementing multi-factor authentication, modernizing legacy systems, or establishing a national CSIRT will provide far greater security benefits than deploying advanced AI technologies. The objective is not to maximize AI adoption, but rather to maximize national resilience.

Integrating AI into National Cyber Priorities

Every national cybersecurity strategy reflects a country's broader national objectives, whether economic development, digital transformation, national security, or public trust. Frontier AI should be evaluated through the same lens.

Governments pursuing digital government initiatives should assess how AI may improve public service delivery while protecting citizen data and maintaining confidence in government systems. Countries seeking to strengthen critical infrastructure should determine whether AI can improve threat detection or operational resilience without introducing unacceptable new risks. Nations facing significant workforce shortages may identify AI-assisted cybersecurity operations or AI-supported training as priority investments.

Importantly, AI should support existing national priorities rather than redefine them. Cybersecurity strategies should remain centered on protecting national interests, enabling economic development, and strengthening resilience, with AI serving as one of several tools available to advance those objectives.

Addressing AI-Specific Risks

As governments evaluate opportunities to use AI, they must also recognize that AI systems introduce new categories of risk that should be reflected within national cybersecurity strategies.[31] [32]

These considerations include securing AI models and training data, protecting sensitive government information used by commercial AI services, managing dependencies on external providers, addressing synthetic media and disinformation, and developing appropriate governance for AI-enabled decision support systems. Governments should also consider how existing cybersecurity regulations, procurement policies, and incident reporting mechanisms apply to AI-enabled technologies.

For many smaller nations, the greatest AI-related risk may not arise from domestically developed frontier models, but from increasing reliance on externally developed AI services from untrusted sources that become embedded throughout government operations. What constitutes a "trusted source" is highly context-dependent and should not be assumed; perceptions of trust in any single provider or country of origin vary considerably across regions and should be assessed independently by each government. This question also intersects with the choice between hosted, closed-weight AI services and open-weight models run on a government's own infrastructure. Running an open-weight model locally can offer greater control over data and infrastructure, but that advantage only holds if the government trusts the provider of the underlying model and can adequately secure its own deployment.

Aligning Cybersecurity and National AI Strategies

[31] https://www.anthropic.com/project/glasswing

[32] https://www.frontiermodelforum.org/technical-reports/managing-advanced-cyber-risks-in-frontier-ai-frameworks/

An increasing number of countries have adopted, or are developing, national AI strategies to promote innovation, economic competitiveness, and responsible AI development.[33] Where these strategies exist, governments should ensure they are closely aligned with national cybersecurity objectives.

The two strategies serve complementary purposes. National AI strategies typically focus on innovation, research, economic development, and responsible AI governance.[34] [35] National cybersecurity strategies focus on protecting digital infrastructure, managing cyber risk, and ensuring national resilience.[36] Together, they should reinforce one another by ensuring that AI systems are secure by design while enabling responsible innovation.

Key Strategic Pillars

Although every country's strategy should reflect its own priorities, several areas deserve particular attention as governments adapt their cybersecurity strategies for the AI era.

Governance and Leadership

Strong governance remains the foundation of effective cybersecurity. Governments should designate clear authorities responsible for cybersecurity policy, AI governance, and implementation while ensuring effective coordination across ministries responsible for technology, defense, intelligence, justice, economic development, and critical infrastructure.

Leadership is particularly important when evaluating AI investments. Governments should establish clear decision-making processes to determine where AI adoption supports national objectives, where additional safeguards are required, and where traditional approaches remain preferable.

Workforce and Skills

AI will not eliminate the need for skilled cybersecurity professionals. Instead, it will change the nature of their work.[37] Governments should invest in developing professionals capable of securely deploying, governing, and evaluating AI-enabled technologies while maintaining strong foundations in cybersecurity, digital governance, and risk management. For many countries, expanding AI literacy across the existing workforce may provide greater benefit than attempting to recruit specialists in frontier AI research. Multi-stakeholder initiatives such as the AI Workforce

[33] https://worldprivacyforum.org/posts/ai-strategies/

[34] https://www.ginc.org/40-national-ai-strategies-august-2025/

[35] https://www.weforum.org/stories/2026/01/five-pathways-shaping-national-ai-investment-strategies/

[36] https://www.centerforcybersecuritypolicy.org/insights-and-research/developing-a-national-cybersecurity-strategy

[37] https://www.weforum.org/stories/2026/05/how-ai-could-open-up-cybersecurity-to-a-wider-workforce/

Consortium are also generating practical guidance on how AI is reshaping ICT and cybersecurity roles, including the competencies and training pathways needed to adapt.[38]

Critical Infrastructure Protection

Critical infrastructure continues to represent one of the highest priorities within any national cybersecurity strategy. Governments should evaluate both how AI may improve resilience and how AI adoption may introduce new operational risks. Infrastructure operators should assess AI-enabled monitoring, predictive maintenance, and automated detection capabilities alongside appropriate governance, human oversight, and cybersecurity protections. AI should strengthen resilience without becoming a single point of failure.[39]

Legacy Systems and Technical Debt

Legacy systems and accumulated technical debt should be treated as a distinct strategic pillar rather than as a secondary consideration of AI adoption. As frontier AI lowers the cost of reconnaissance and exploitation, unmodernized systems, including those that cannot support multi-factor authentication, rapid patching, or modern network segmentation, become disproportionately exposed.[40] Reducing this attack surface is one of the most cost-effective steps governments can take to improve national cyber resilience in the AI era.

This priority is reinforced by recent guidance from the Five Eyes cybersecurity agencies and by supervisory expectations following the European Central Bank's letter to banks, both of which call on governments and regulated entities to accelerate vulnerability management and patching, reduce attack surface, and secure open-source software libraries without delay. Governments should incorporate legacy system modernization into national cybersecurity strategies as an immediate priority, alongside, and in many cases ahead of, new investments in frontier AI.[41]

Technology and Innovation

Governments should encourage innovation while ensuring that technology adoption aligns with national priorities and available resources. Frontier AI is only one component of a broader digital ecosystem that includes secure digital infrastructure, cloud computing, identity management, and resilient cybersecurity capabilities. For many smaller nations, strengthening these foundations will provide greater near-term benefits than pursuing frontier AI alone.

Rather than attempting to develop or host frontier AI models, most governments will achieve greater value by leveraging commercially available AI-enabled cybersecurity tools and services.

---

[38] https://www.cisco.com/site/m/ai-workforce-consortium/index.html

[39] https://digiamericas.org/wp-content/uploads/2026/05/SC_EN.pdf

[40] https://blogs.cisco.com/news/scaling-ai-in-the-enterprise-how-technical-debt-limits-returns-on-ai

[41] https://www.nsa.gov/Press-Room/News-Highlights/Article/Article/4523810/five-eyes-cyber-security-agencies-statement/

Many cybersecurity functions, such as threat intelligence summarization, policy development, log analysis, and security operations, can be effectively supported using commercially available or smaller AI models. More advanced frontier models may provide additional value for specialized activities such as dynamic analysis of compiled software, advanced vulnerability research, or complex cyber investigations. Governments should therefore match AI capabilities to operational requirements rather than assume the most advanced model is always necessary. Regardless of AI adoption decisions, governments should continue prioritizing established cybersecurity practices such as identity security, asset management, vulnerability management, logging, monitoring, and network segmentation.

Governments should also focus on building organizational AI maturity by integrating AI into existing cybersecurity workflows, developing governance processes, and training personnel to effectively use AI-assisted tools. Wherever possible, they should adopt model-agnostic approaches that allow organizations to evaluate and integrate different AI technologies as capabilities, costs, and availability evolve. This flexibility reduces dependence on any single provider while positioning governments to take advantage of future advances in frontier AI.[42]

Governments should avoid optimizing their cybersecurity strategies around the capabilities of any single AI model. Instead, they should build flexible processes that allow new models to be evaluated and integrated as capabilities, costs, and availability evolve. Ultimately, technology investments should be guided by the same principle as every other cybersecurity investment: adopt AI where it demonstrably improves cybersecurity outcomes and represents the best use of limited national resources. Purpose-built security harnesses and orchestration layers built around underlying models are often as important as the models themselves, since they translate raw model capability into reliable, auditable security outcomes.[43]

International Cooperation

Because relatively few countries will possess the resources necessary to develop frontier AI independently, international cooperation becomes increasingly important. Governments should continue participating in regional cybersecurity initiatives, international standards development, information sharing organizations, and capacity-building programs while exploring opportunities for trusted AI partnerships. Shared expertise, regional SOCs, joint training initiatives, and collaborative research may allow smaller nations to benefit from frontier AI without independently bearing its full costs.

A Whole-of-Society Approach

[42] https://blogs.cisco.com/ai/announcing-foundry-security-spec

[43] https://blogs.cisco.com/news/8-years-of-security-research-in-8-weeks-transforming-cybersecurity-with-ai

Cybersecurity cannot be achieved through government action alone. That principle is even more important in the AI era.

Government

Government provides national leadership by establishing policy, coordinating implementation, protecting public systems, and managing national cyber risk. It should also ensure that AI adoption within government reflects principles of transparency, accountability, security, and public trust.

Private Sector

The private sector owns and operates much of the infrastructure that AI will transform. Technology providers, telecommunications companies, cloud service providers, cybersecurity firms, and critical infrastructure operators all play essential roles in identifying practical AI use cases, sharing threat intelligence, and implementing security measures. The private sector also holds much of the responsibility for addressing legacy systems and accumulated technical debt within critical infrastructure, both of which increase vulnerability to AI-enabled attacks and should be prioritized alongside new AI adoption. Governments should engage industry early when considering AI-related policies to ensure that proposed requirements are technically feasible, economically sustainable, and aligned with international standards.

Academia

Universities and research institutions contribute to workforce development, technical expertise, and independent analysis of emerging AI technologies. Academic institutions can also support governments by evaluating AI risks, conducting applied cybersecurity research, and developing educational programs that prepare future cybersecurity professionals.

Civil Society

Public trust remains central to successful cybersecurity policy. Civil society organizations contribute expertise on privacy, human rights, digital inclusion, consumer protection, and public education. Their participation helps ensure that AI-related cybersecurity policies reflect national values while maintaining public confidence in government institutions.

Ultimately, the same collaborative approach that has long characterized effective national cybersecurity strategies remains the best model for the AI era. Frontier AI changes the technologies available to governments, but it does not change the fundamental requirement for coordinated leadership, shared responsibility, and sustained investment across society.

The remaining question is not whether AI should appear in a national cybersecurity strategy, but how governments can use AI, where appropriate, to achieve the strategy's objectives. The

following section examines how AI can support implementation by strengthening cyber defenses, expanding national capacity, and improving the day-to-day execution of cybersecurity programs.

4. Using AI to Implement a National Cybersecurity Strategy

A national cybersecurity strategy succeeds or fails not because of the quality of the document itself, but because of its implementation. Establishing strategic objectives, governance structures, and policy priorities is only the beginning. Governments must translate those objectives into operational capabilities, institutional responsibilities, and measurable outcomes.

The Center's previous work emphasized that implementation requires sustained investment in governance, workforce development, critical infrastructure protection, information sharing, incident response, and continuous evaluation.[44] These remain the core building blocks of an effective national cybersecurity program. Frontier AI does not replace these capabilities, nor does it eliminate the need for skilled professionals, sound policy, or institutional leadership. Instead, where appropriate, AI can help governments execute these functions more efficiently, particularly where resources are limited and analytical workloads continue to grow.

Governments should develop model-agnostic AI workflows that allow organizations to evaluate multiple AI models as capabilities, costs, and availability evolve. This reduces dependence on any single provider and allows organizations to match the complexity of the model to the task being performed.[45]

The question for policymakers is therefore not whether AI should be integrated into every aspect of cybersecurity implementation, but where it provides measurable operational value that justifies the investment.

Strengthening National Cyber Defenses

National Cybersecurity Centers

Many countries have established centralized cybersecurity organizations like National Cybersecurity Centers, Computer Security Incident Response Teams (CSIRTs), or Computer Emergency Response Teams (CERTs) to coordinate national cyber defense, provide technical guidance, share threat intelligence, and support incident response. For smaller nations, these organizations often serve as the operational hub of national cybersecurity efforts.

[44]https://www.centerforcybersecuritypolicy.org/insights-and-research/developing-a-national-cybersecurity-strategy
[45]https://blogs.cisco.com/ai/announcing-foundry-security-spec

Governments should first ensure that these organizations possess clearly defined authorities, sustainable funding, and strong relationships across government and the private sector. Investments in governance, staffing, and operational maturity are generally more important than investments in advanced technologies.[46]

Where operational needs justify it, AI may enhance these organizations by assisting analysts with processing large volumes of security data, correlating threat intelligence from multiple sources, summarizing technical reporting, identifying emerging trends, and prioritizing incidents requiring immediate attention.[47] These capabilities should be viewed as tools that augment human analysts rather than replace them. Human expertise remains essential for contextual analysis, operational decision-making, and coordination during significant cyber incidents.[48] Countries without mature national cyber centers should generally prioritize establishing these foundational capabilities before pursuing more advanced AI-enabled operational functions.

Building AI Workflows Before Frontier AI

Governments do not need access to frontier AI models to begin realizing the benefits of AI in cybersecurity. Many routine functions, including threat intelligence analysis, documentation, malware triage, and policy research, can already be enhanced using commercially available AI models and AI-enabled security products. By first integrating AI into these existing workflows, organizations can develop the operational experience needed to evaluate whether more advanced frontier AI capabilities provide sufficient additional value to justify their higher costs and complexity.

Incident Response

Rapid detection and response remain among the most effective methods for limiting the impact of cyber incidents. Governments should ensure that national incident response procedures are clearly defined, regularly exercised, and supported by strong coordination among relevant agencies and critical infrastructure operators. AI may improve portions of the incident response process by assisting with alert triage, malware analysis, forensic investigation, vulnerability prioritization, and documentation.[49] These capabilities may reduce the time required to understand an incident and recommend appropriate remediation actions.

However, governments should avoid over-automating incident response. Significant cyber incidents often require policy judgments, legal considerations, diplomatic coordination, and public

---

[46] https://www.oecd.org/en/publications/digital-government-review-of-korea_9defc197-en/full-report/strengthening-governance-investment-and-skills-for-digital-government_a4dabc49.html

[47] https://atlas.mitre.org/

[48] https://www.isaca.org/resources/isaca-journal/issues/2021/volume-5/what-role-do-humans-play-in-ensuring-cybersecurity

[49] https://www.anthropic.com/research/glasswing-initial-update

communication that extend well beyond technical analysis. AI should support—not replace—human decision-makers during these events.

Threat Intelligence

Threat intelligence enables governments to anticipate emerging risks rather than merely react to cyber incidents after they occur. Effective national cybersecurity programs integrate technical intelligence with geopolitical developments, criminal trends, vulnerability disclosures, and information shared by international partners.

One of AI's most practical applications may be assisting analysts in processing the enormous volume of cybersecurity reporting generated each day. AI can summarize reports, identify common indicators, translate foreign-language reporting, organize technical information, and highlight trends relevant to national priorities. Used appropriately, these capabilities allow analysts to spend less time processing information and more time assessing its implications for national security.

The effectiveness of threat intelligence ultimately depends not on the sophistication of analytical tools, but on trusted information sharing relationships among government agencies, the private sector, international partners, and the cybersecurity community.

Expanding National Cyber Capacity

Workforce Development

No technology can compensate for the absence of a capable cybersecurity workforce. The Center's previous paper identified workforce development as one of the highest-return investments any nation can make, and that conclusion remains unchanged. Resources such as the AI Workforce Consortium, which is publishing research specifically addressing how AI is changing the skills and roles required of cybersecurity professionals, can help inform these investments.[50]

Governments should continue investing in cybersecurity education, professional development, technical training, certification programs, and career pathways that strengthen national cyber capacity over the long term.

AI may assist these efforts by providing personalized learning, interactive training environments, simulated cyber exercises, automated tutoring, and support for cybersecurity professionals learning new technologies. Countries facing persistent workforce shortages may also find that AI enables experienced analysts to supervise larger teams or reduces time spent on repetitive analytical tasks.

[50] https://www.cisco.com/site/m/ai-workforce-consortium/index.html

Nevertheless, AI should not be viewed as a substitute for workforce development. Skilled professionals remain essential for governance, policy development, risk management, and operational decision-making.

Public Awareness

A resilient nation depends not only on government capabilities but also on informed citizens, businesses, and public institutions. Public awareness campaigns remain one of the most cost-effective methods for reducing cyber risk by improving digital hygiene, encouraging incident reporting, and reducing susceptibility to fraud and social engineering.

As AI-generated phishing, synthetic media, and online impersonation become increasingly sophisticated, public education should evolve accordingly. Governments should ensure that citizens understand how AI may be used to facilitate fraud, misinformation, and identity theft while providing practical guidance for identifying suspicious communications and verifying trusted information sources.

AI itself may support public awareness initiatives through multilingual communication, personalized educational materials, and improved accessibility. However, governments should continue emphasizing trusted human communication, particularly during emergencies where public confidence is essential.

Government Training

Governments increasingly rely on digital technologies across every ministry and agency.[51] Consequently, cybersecurity cannot remain the responsibility of technical personnel alone. Senior leaders, policymakers, procurement officials, legal advisors, communications professionals, and operational managers should all receive appropriate cybersecurity education relevant to their responsibilities. As AI becomes more widely adopted across government, this training should include both opportunities and risks associated with AI-enabled systems.

Governments considering AI adoption should also ensure that officials understand issues such as data protection, responsible AI use, procurement considerations, model limitations, and appropriate human oversight.

Developing Organizational AI Maturity

Successfully integrating frontier AI into cybersecurity is as much an organizational challenge as it is a technical one. Before pursuing the most advanced AI capabilities, governments should focus on building experience using AI within existing cybersecurity workflows. This includes identifying

[51] https://www.oecd.org/en/publications/digital-government-outlook_0496b2bc-en/full-report.html

appropriate use cases, establishing governance and oversight, training personnel, and developing processes for validating AI-generated outputs.

For many governments, commercially available AI models, open-weight models, and AI-enabled cybersecurity products provide an effective starting point. These technologies allow organizations to gain practical experience while building the institutional knowledge needed to evaluate and securely adopt more capable frontier AI systems as they become more accessible and cost-effective.

Ultimately, successful AI adoption depends on more than access to powerful models. It requires mature cybersecurity processes, secure integration, appropriate testing and evaluation methodologies, and a skilled workforce capable of understanding both the capabilities and limitations of AI. Governments that invest in building these organizational capabilities today will be better prepared to realize the benefits of frontier AI in the future.

Cybersecurity Fundamentals Remain Essential

Frontier AI has the potential to improve many aspects of cybersecurity, but it does not reduce the importance of established cybersecurity fundamentals. In fact, AI creates a step change in their importance and accelerates the need for governments to act. Most successful cyber intrusions continue to exploit well-known weaknesses such as compromised credentials, excessive privileges, insecure configurations, unpatched systems, and inadequate network visibility. Legacy systems are particularly exposed in this environment: they are harder to observe, easier for autonomous AI tools to probe, and often cannot support modern security protections such as multi-factor authentication or rapid patching. Many also rely on non-memory-safe programming languages that the technology industry is actively working to phase out through code refactoring, leaving unmodernized systems increasingly exposed to the classes of vulnerabilities such efforts are designed to eliminate. While frontier AI may enable organizations to identify these issues more quickly or prioritize them more effectively, it does not eliminate the need to address them.

Governments should therefore continue investing in proven cybersecurity practices as the foundation of national resilience. These include implementing multi-factor authentication (MFA), adopting Zero Trust Architecture (ZTA) principles, enforcing least-privilege access, maintaining accurate asset inventories, applying timely security patches, segmenting critical networks, and ensuring comprehensive logging, monitoring, and vulnerability management. These measures remain among the most effective and cost-efficient ways to reduce cyber risk.

For many governments, particularly those with constrained cybersecurity resources, strengthening these foundational capabilities will provide greater near-term improvements to national cybersecurity than investing in advanced AI technologies alone. Frontier AI should be viewed as a force multiplier that enhances existing cybersecurity programs as opposed to a

substitute for the governance, processes, and technical controls that have long formed the basis of effective cyber defense. Countries that first establish strong cybersecurity fundamentals will be better positioned to realize the benefits of frontier AI as these technologies continue to mature.

Supporting Critical Infrastructure

Risk Management

Protecting critical infrastructure remains a cornerstone of every national cybersecurity strategy. Governments should continue encouraging risk-based security practices that reflect sector-specific threats and operational realities. Risk management in the AI era should also explicitly consider how AI is reshaping the broader threat landscape and, in turn, how it affects an organization's ability to identify, protect, detect, respond, and recover from cyber incidents, not solely the risks associated with adopting AI tools themselves.[52]

As operators evaluate AI-enabled technologies, they should assess whether these tools meaningfully improve resilience relative to their costs, technical complexity, and associated risks. In some sectors, AI-assisted anomaly detection or predictive maintenance may produce measurable operational benefits. In others, simpler investments such as network segmentation, asset visibility, or vulnerability management may provide greater improvements to security.

Governments should avoid viewing AI adoption as a prerequisite for resilient infrastructure. Rather, AI should be evaluated alongside other investments using the same principles of risk management and cost-effectiveness.

Resilience

Resilience requires that essential services continue functioning despite cyber disruptions.[53] Governments should encourage regular cyber exercises, continuity planning, recovery procedures, and coordination mechanisms across sectors.

As organizations increasingly rely on AI-enabled systems, resilience planning should also account for the possibility that these systems themselves become unavailable, compromised, or produce unreliable outputs. Human operators should retain the ability to validate critical decisions and maintain essential services if AI systems fail or cannot be trusted. Ultimately, resilient infrastructure depends more on sound governance, redundancy, and preparedness than on any single technology.

Public-Private Collaboration

[52]https://www.cisa.gov/cross-sector-cybersecurity-performance-goals/cross-sector-cybersecurity-performance-goals
[53]https://www.cisco.com/site/us/en/learn/topics/security/what-is-cyber-resilience.html

Because much of a nation's digital infrastructure is owned or operated by the private sector, successful implementation depends on sustained collaboration between government and industry.

Governments should continue developing trusted mechanisms for information sharing, joint exercises, coordinated incident response, and collaborative policy development. Industry should be engaged early when governments consider AI-related regulations, procurement requirements, or cybersecurity standards to ensure that policies are practical, technically feasible, and aligned with international best practices.

Technology providers also possess valuable operational experience regarding AI deployment, security, and risk management. Leveraging this expertise through structured partnerships may prove more valuable for many smaller nations than attempting to independently develop frontier AI capabilities.

Measuring Progress

AI-Enabled Monitoring and Assessment

Implementation requires continuous measurement. Governments should establish clear metrics to evaluate whether strategic objectives are being achieved, whether cybersecurity investments are improving resilience, and whether resources are being allocated effectively.

AI may assist by analyzing operational data, identifying trends across agencies, summarizing implementation progress, and supporting periodic reviews. These capabilities can improve visibility into national cybersecurity programs while reducing administrative burdens associated with reporting and performance measurement.

However, AI-generated assessments should complement human evaluation rather than replacing it. Strategic decisions require consideration of political, economic, legal, and societal factors that extend beyond technical metrics.

Strategy Review and Updates

A national cybersecurity strategy should be treated as a living document that evolves alongside technology, threats, and national priorities. Governments should establish regular review cycles that evaluate changes in the threat landscape, emerging technologies, workforce needs, implementation progress, and international developments. These reviews should also assess whether previous assumptions regarding AI remain valid. Advances in AI capabilities, reductions in deployment costs, or new regional partnerships may create opportunities that were not feasible when the strategy was originally developed.

Equally important, governments should periodically reassess whether AI investments continue to represent the best use of limited national resources. The rapid pace of AI development means that today's strategic decisions may require adjustment as technologies mature and national priorities evolve.

Ultimately, successful implementation is not measured by the extent to which a government adopts artificial intelligence. It is measured by whether the nation becomes more secure, more resilient, and better prepared to manage cyber risk. For some countries, AI will become a significant force multiplier across multiple aspects of cybersecurity implementation. For others, its role may remain limited while foundational capabilities continue to mature. In either case, implementation should remain guided by evidence, national priorities, and a realistic assessment of available resources rather than by technological enthusiasm alone.

5. Decision Roadmap for Governments with Constrained Cybersecurity Resources

No single roadmap will be appropriate for every country. Countries differ considerably in their digital maturity, cybersecurity workforce, financial resources, and national priorities. Some governments may already operate mature national cyber centers and digital government platforms, while others are still establishing foundational cybersecurity capabilities. Frontier AI should therefore be viewed as one of many strategic tools available to governments rather than a universal solution.

The following roadmap provides a practical framework for evaluating where frontier AI may contribute to national cybersecurity objectives while recognizing that countries should adopt these capabilities at a pace consistent with their own circumstances.

Immediate Actions (0–12 Months)

The first priority for governments should be understanding how frontier AI affects their national cybersecurity posture rather than immediately investing in AI technologies.

Governments should designate clear leadership responsible for evaluating the cybersecurity implications of frontier AI and coordinating policy across relevant ministries and agencies. This is an area where the window to respond is narrow: joint guidance from a multinational group of cybersecurity agencies underscores the urgency of acting now rather than treating this as a longer-term priority.[54] As part of this effort, national risk assessments should be updated to

[54]https://www.nsa.gov/Press-Room/News-Highlights/Article/Article/4523810/five-eyes-cyber-security-agencies-statement/

consider AI-enabled cyber threats, dependencies on external AI providers, and opportunities for AI to strengthen existing cybersecurity capabilities.

Rather than launching large-scale AI programs, governments should identify a small number of pilot projects that address clearly defined operational challenges. Examples may include using commercial frontier AI models to assist with policy analysis, summarize threat intelligence, support cybersecurity awareness campaigns, or improve internal government workflows. Governments should also establish baseline metrics that allow them to measure whether AI improves cybersecurity outcomes relative to existing approaches. These limited deployments allow governments to evaluate costs, operational benefits, governance requirements, and security considerations before making larger investments.

Medium-Term Priorities (1–3 Years)

As governments gain experience, they should determine whether frontier AI can strengthen existing cybersecurity capabilities in areas where measurable operational benefits exist.

Potential applications include supporting national cyber centers with threat intelligence analysis, assisting incident response teams with investigation and documentation, improving cyber workforce training, or enhancing vulnerability management. Governments should continue investing in the foundational capabilities identified throughout this paper, including governance, workforce development, information sharing, and critical infrastructure protection, while incorporating frontier AI where it demonstrably improves operational effectiveness.

During this period, governments should also establish governance frameworks that address the secure adoption of frontier AI, including procurement policies, data protection requirements, risk management processes, human oversight, and coordination with existing national AI strategies where appropriate.

Long-Term Vision (3–5 Years)

Over the longer term, governments should seek to build sustainable national cyber resilience rather than simply increasing AI adoption. For some nations, frontier AI may become an integral component of cybersecurity operations, supporting analysts, accelerating threat detection, improving policy development, and strengthening decision-making across government. For others, the greatest long-term value may come through trusted partnerships with commercial providers, regional cybersecurity organizations, or international allies that provide access to frontier AI capabilities without requiring significant domestic investment.

Regional cooperation is likely to become increasingly important. Shared Security Operations Centers, regional threat intelligence exchanges, collaborative workforce development initiatives,

and joint research partnerships can help smaller nations benefit from frontier AI while reducing costs and avoiding unnecessary duplication of effort.

Ultimately, the objective is not to build sovereign frontier AI capabilities for their own sake, but to establish sustainable cybersecurity institutions capable of adapting as technologies, threats, and national priorities continue to evolve.

6. Conclusion

Frontier AI is rapidly reshaping the cybersecurity landscape. It is changing how cyber threats are developed, how defenders respond, and how governments must think about national resilience. For smaller nations, these changes create both new risks and new opportunities. Cyber adversaries are increasingly leveraging frontier AI to automate attacks, accelerate vulnerability discovery, and enhance social engineering at a scale that was previously unattainable. At the same time, the same technologies have the potential to augment limited cyber workforces, improve threat intelligence, strengthen incident response, and make cybersecurity capabilities more accessible than ever before.

The implications extend beyond technology alone: frontier AI is changing the economics of cybersecurity. Capabilities that once required large budgets, specialized expertise, and years of institutional development are becoming increasingly available through commercial AI services and AI-enabled security products. This evolution presents smaller nations with an opportunity to strengthen their cybersecurity posture without necessarily replicating the extensive cyber institutions of larger states. However, realizing these benefits requires deliberate planning, sound governance, and a clear understanding of where frontier AI creates measurable value.

Not every country should pursue the same path. Smaller nations differ significantly in their digital maturity, available financial resources, cybersecurity workforce, infrastructure, and strategic priorities. For some governments, frontier AI may become a force multiplier that enhances national cyber operations, improves policy development, and enables more effective use of limited resources. For others, the most effective strategy may be to continue investing in foundational cybersecurity capabilities while selectively adopting frontier AI through trusted commercial providers or regional partnerships. Neither approach is inherently correct or incorrect. What matters is that governments make these decisions intentionally, based on evidence, national priorities, and a realistic assessment of available resources.

The principles for developing and implementing a national cybersecurity strategy remain unchanged. Effective strategies begin with understanding national risks, establishing clear objectives, engaging stakeholders, building strong governance, investing in workforce development, protecting critical infrastructure, and continuously evaluating progress. Frontier AI should not replace these principles; it should reinforce them. Governments that view AI as one

component of a broader cybersecurity strategy rather than an end in itself will be better positioned to make sustainable investments that strengthen national resilience over time.

The most important question facing policymakers is not whether they should adopt frontier AI, but how frontier AI changes the decisions they must make. Every government should assess how AI is changing its threat environment, determine where AI can meaningfully improve cybersecurity outcomes, evaluate whether those benefits justify the associated costs and risks, and continually reassess those decisions as both technologies and national priorities evolve. This process of continuous evaluation is essential because the pace of AI development is unlikely to slow, and strategies that remain static will quickly become outdated.

Ultimately, success should not be measured by the extent to which a nation adopts frontier AI. It should be measured by whether the nation becomes more secure, more resilient, and better prepared to manage cyber risk. Frontier AI is a powerful new capability, but it is only one element of national cybersecurity. Leadership, governance, trusted partnerships, technical expertise, and sustained investment remain the foundations of cyber resilience.

For smaller nations, the future is not a choice between embracing frontier AI or ignoring it. Frontier AI will influence the cybersecurity environment regardless of whether governments actively deploy it. The strategic opportunity lies in understanding where it can advance national objectives, where it introduces new risks, and where traditional investments continue to provide greater value. Countries that approach frontier AI with the same discipline, pragmatism, and strategic planning that characterize effective national cybersecurity strategies will be best positioned to protect their citizens, strengthen their economies, and build resilient digital societies in an increasingly AI-enabled world.