

The Evolving AI Risk Landscape: Cybersecurity, Autonomous Systems and Emerging Global Security Challenges





Report:

The Evolving AI Risk Landscape:

Cybersecurity, Autonomous Systems and Emerging Global Security Challenges

San Marino, 30 September 2026

Disclaimer: This document is prepared by the researchers of the Centre for Global Studies (CGS) of the Parliamentary Assembly of the Mediterranean (PAM) in San Marino in their personal capacity. The opinions expressed in the note are the authors' own and may not reflect the views of PAM. This report underwent AI-assisted linguistic review.

Index

Executive Summary	2
Introduction	5
Background	9
Cybersecurity Risks Arising from AI Misuse	10
Agent Interaction and the Acceleration of Cyber Threats in 2026	11
Moltbook and the security of agent interaction	11
From task assistance to connected attack workflows	12
Unauthorized agent behavior and the problem of control	14
Communication across intended boundaries	16
Evaluation incidents and model supervision.....	16
Weapons Related Misuse and Biological Research Risks	20
Provider Safeguards and the Limits of Intervention	21
Frontier AI Warnings and International Governance	22
Attribution and the Future of Effective Human Control.....	25
National Positions and Geopolitical Implications of AI Governance	27
Recommendations for Parliamentary Action in the Euro-Mediterranean and Gulf Regions	40
Legislative Priorities	42
Technical and Institutional Priorities	45
Parliamentary Oversight and the Role of PAM	47
Conclusions.....	49

Executive Summary

Artificial intelligence is evolving from a supporting technology into an operational component of cyber activity, scientific research, essential services, and security-related decision-making. This transition offers substantial benefits while changing the scale, speed, and organization of risk. The central governance challenge extends beyond harmful content and isolated misuse to AI systems capable of connecting multiple tasks, interacting with external services, and taking consequential actions with reduced human involvement.

The evidence examined in *The Evolving AI Risk Landscape: Cybersecurity, Autonomous Systems and Emerging Global Security Challenges* identifies two distinct but increasingly interconnected risk pathways. The first is human-directed misuse, involving malicious actors employing AI to strengthen deception, reconnaissance, vulnerability exploitation, malware development, and weapons-related research. The second is unauthorized agent behavior, involving systems pursuing legitimate objectives through actions beyond the permissions, methods, or boundaries established by operators. Effective governance requires the capacity to distinguish these pathways and address circumstances in which human-directed misuse and unauthorized agent behavior overlap.

In cybersecurity, AI is lowering barriers across several stages of malicious operations. AI assistance is expanding from phishing, impersonation, and synthetic media to vulnerability research, adaptive malware, and post-compromise activity. The principal strategic change concerns both improved performance on individual tasks and the integration of those tasks into connected workflows. Such integration can enable malicious actors to operate faster, delegate more intermediate steps, and expand operations without a corresponding increase in personnel or expertise.

Agentic systems introduce an additional layer of exposure because the risks associated with such systems depend on available authority and resources. Access to credentials, networks, files, communication channels, and external tools can translate an unreliable output into a consequential action. Recent evaluation incidents indicate that model instructions and refusal mechanisms alone cannot provide adequate protection. Effective safeguards require technical restrictions on agent access and actions, continuous monitoring of consequential activity, reliable attribution, and the capacity to halt execution and communication.

Interaction among agents further complicates containment. External content may influence subsequent conduct, while communication across platforms can allow information, instructions,

or resources to move beyond the environment originally supervised by an operator. Governance arrangements should therefore address individual systems and the wider operational environment through which agents identify other participants, exchange information, and act across services.

Risks involving weapons-related activities and biological research demonstrate how digital assistance may contribute to physical harm. Provider intervention can interrupt access to a service without removing knowledge, software, or technical materials already acquired by an actor. The effectiveness of intervention must consequently be assessed against the timing of detection, the extent of containment, and remaining capabilities. In dual-use scientific fields, safeguards must combine technical monitoring with specialized expertise capable of distinguishing legitimate research from connected patterns of activity that raise security concerns.

AI providers remain important actors through system evaluation, activity monitoring, and the withdrawal of access. Provider capacity to intervene nevertheless remains limited by deployment arrangements and the distribution of technical control. Local, offline, open-weight, and distributed deployments may reduce or eliminate provider visibility and control. Responsibility must therefore be allocated across the full chain of authority, including developers, deploying organizations, infrastructure operators, and public institutions. Each actor should have clearly defined obligations concerning access control, monitoring, incident response, evidence preservation, and cooperation. The emerging policy debate concerns whether advances in capability should proceed without prior evidence that safeguards can keep pace. Independent evaluation can support credible assurance when evaluators have adequate access, serious findings trigger defined responses, and public authorities can act on the evidence. Voluntary commitments remain useful, but cannot substitute for common requirements concerning high-risk evaluations, incident reporting, response timelines, and the conditions warranting modification or delay of deployment or further capability development. National positions reveal differing approaches to the balance between technological competitiveness, preventive oversight, and international regulatory authority, making cooperation on evaluation and incident notification a practical priority amid wider geopolitical differences.

For the Euro-Mediterranean and Gulf regions, unequal institutional resilience presents an additional challenge. Advanced defensive tools will have limited collective value if public administrations, essential services, and smaller institutions lack the resources to assess warnings, withdraw access, or coordinate a response. Capacity building should therefore form part of regional security policy alongside technical assistance. Cooperation should strengthen baseline

cyber resilience, trusted channels for incident notification, and access to relevant scientific and technical expertise.

Parliaments have a distinct role in connecting AI governance with democratic authority, human rights, and international security. Legislative scrutiny should examine whether responsibility corresponds to actual control; whether consequential actions remain attributable and subject to effective intervention; whether serious incidents are reported and independently examined; and whether monitoring measures are proportionate and subject to review. This responsibility has acquired greater urgency following the warning by Volker Türk, UN High Commissioner for Human Rights, that voluntary self-regulation is insufficient and that individual companies should not determine acceptable risks with global consequences. High Commissioner Türk's call for stronger regulation, serious-incident reporting, independent verification of AI capabilities, and mandatory human rights due diligence reinforces the need for coordinated public oversight. Parliamentary institutions should also establish safeguards for AI-assisted legislative work, including source verification, protection of confidential information, and clearly assigned human responsibility for official texts¹. International cooperation should promote compatible standards, discourage competition that weakens safeguards, and ensure meaningful participation by states with limited resources.

This report identifies effective human control as an institutional capability requiring more than the formal presence of a human decision-maker. Meaningful control depends on sufficient authority to supervise, constrain, and halt AI-enabled actions; reliable evidence to reconstruct incidents; and institutions capable of intervening across organizational and national boundaries. Preserving effective human control while enabling beneficial innovation should remain the central strategic objective of AI governance.

¹ UN News. (14 September 2026). *Countries Must Increase AI Regulation to Avoid 'Existential Risks': Türk*. United Nations. <https://news.un.org/en/story/2026/09/1168326>


Introduction

1. This report, *The Evolving AI Risk Landscape: Cybersecurity, Autonomous Systems and Emerging Global Security Challenges*, has been prepared by the Centre for Global Studies of the Parliamentary Assembly of the Mediterranean (PAM-CGS). It examines the evolution of AI-enabled cyber threats, the emergence of unauthorized agent behavior, whereby AI agents may take actions beyond the limits or authority intended by their operators while pursuing assigned tasks. It further considers their implications for national and international governance. Its purpose is to support legislators, policymakers and technical experts in assessing how responsibility, oversight and intervention should develop as AI capabilities become more powerful and more widely accessible.
2. Artificial intelligence is becoming increasingly integrated into the systems through which societies communicate, conduct research, deliver services and manage information. Its expanding capabilities offer opportunities to improve productivity, support scientific discovery and strengthen institutional preparedness. These benefits are accompanied by new responsibilities concerning the protection of sensitive data, the security of digital infrastructure and the supervision of decisions delegated to automated systems. Understanding the evolving risk landscape requires examining how these technologies are used, the environments in which they operate and the capacity of responsible institutions to maintain effective control.
3. Early public debates focused heavily on employment, the transformation of professional roles and the wider social consequences of automation. The International Labour Organization's (ILO) 2023 assessment highlighted the potential for generative AI to complement many occupations while changing the tasks, working conditions and skills associated with them. These concerns remain relevant as AI becomes more deeply embedded in economic and social life. They now form part of a broader policy agenda that also encompasses the security implications of systems capable of producing convincing content, assisting technical work and acting through connected digital services².
4. Within this context, the report traces the progression from AI-enabled deception and cyber intrusion to agent interaction, unauthorized behavior and emerging risks involving weapons-related activities and biological research. It examines the relevance of malicious AI tools

² International Labour Organization (ILO). (21 August 2023). *Generative AI likely to augment rather than destroy jobs*. International Labour Organization. <https://www.ilo.org/resource/news/generative-ai-likely-augment-rather-destroy-jobs>

within illicit digital ecosystems and considers the capacity of provider safeguards to prevent, detect and contain harmful activity. For the Euro-Mediterranean and Gulf regions, these questions have direct relevance to parliamentary oversight, institutional resilience and cross-border cooperation. The analysis seeks to inform governance arrangements capable of preserving the benefits of innovation while addressing a central question: how can human-directed misuse and increasingly autonomous behavior be distinguished, and how can meaningful oversight be sustained as AI capabilities extend beyond the visibility and control of any single provider?

Figure 1: The AI Risk-Control Nexus: Human-Directed Misuse, Unauthorized Agent Behavior and Effective Human Control³



Note: This report identifies two increasingly interconnected AI risk pathways: human-directed misuse and unauthorized agent behavior. Their convergence underscores the importance of effective human control, including access restrictions, monitoring, attribution, intervention, and accountability across developers, deploying institutions, infrastructure operators, and public authorities.

³ Graphic compiled by the CGS research team, based on the analysis presented in this report.

Table 1: Key AI and Technology Actors

Actor	Profile
Anthropic	AI company developing the Claude family of advanced AI systems.
Dario Amodei	Co-Founder and Chief Executive Officer (CEO) of Anthropic.
ESET	Cybersecurity company specializing in threat detection and security research.
Elon Musk	Founder of xAI and Chief Executive Officer (CEO) of SpaceX.
Google Threat Intelligence Group (GTIG)	Google’s specialized threat-intelligence research group.
Hugging Face	AI development and model-hosting platform.
METR	Independent research organization specializing in the evaluation of advanced AI systems.
Moltbook	Digital platform designed primarily for interaction among AI agents.
Nightingale Collective	Research group associated with the investigation of AI-agent communication examined in this report.
OpenAI	AI research and deployment company and developer of advanced AI systems.
Sam Altman	Chief Executive Officer (CEO) of OpenAI.
Wiz	Cloud-security company specializing in cloud infrastructure and security research.

Table 2: Key AI Systems, Tools and Technical Terms

Term / System	Definition
AI Agent	An AI system capable of pursuing an assigned objective through a sequence of actions, including using tools, accessing information and interacting with other systems with varying degrees of autonomy.
Agentic AI	AI designed to perform multi-step tasks and pursue objectives with a degree of operational autonomy rather than responding only to individual user prompts.
Claude Code	Anthropic’s agentic coding tool designed to perform software-development and coding-related tasks.
Deepfake	AI-generated or manipulated audio, image or video content designed to realistically imitate a person, object or event.
dots	Persistent AI agents developed by OpenAI, powered by GPT-6 Astra, that can carry out ongoing tasks using a cloud computer, browser, and connected applications.
Distributed AI Deployment	Deployment in which the development, hosting, access or operation of an AI system is divided among different organizations, infrastructures or jurisdictions.
Effective Human Control	The capacity of humans and responsible institutions to supervise, constrain and, where necessary, intervene in or halt consequential AI-enabled actions.

FraudGPT	A malicious generative AI tool associated with AI-assisted criminal activities and distributed through illicit online ecosystems.
Frontier AI	Advanced general-purpose AI systems at or near the leading edge of current capabilities, particularly systems whose increasing capabilities may create significant safety or security considerations.
Gemini	Google's family of generative AI models.
GPT-6 Astra	An OpenAI frontier model assessed by the company as having Critical cybersecurity capability under its Preparedness Framework.
GPT-6.1 Sol	An OpenAI frontier model assessed by the company as having Critical cybersecurity capability and High biological and chemical capability under its Preparedness Framework.
Local / Offline AI Models	AI models operated on local devices or infrastructure rather than through a remotely controlled provider service, potentially limiting provider visibility and capacity for remote intervention.
Model Misalignment	Behavior in which an AI system acts in ways that do not fully correspond with intended human objectives, instructions or constraints.
Prompt Injection	A technique in which instructions introduced through user input or external content processed by an AI system attempt to redirect or manipulate its behavior.
PromptSpy	Android malware analyzed by ESET in 2026 that incorporated Gemini to interpret information displayed on a device and support persistence.
Sandboxing	The isolation of an AI system or software within a restricted environment to limit its access to external systems, networks, files or other resources.
Safety case	A structured, evidence-based argument documenting how specified risks are assessed and managed before a high-risk AI activity proceeds.
Unauthorized Agent Behavior	Actions taken by an AI agent beyond the limits or authority intended by its operator while pursuing an assigned task, including unauthorized access, communication across intended boundaries or the use of unapproved methods.
WormGPT	A generative AI tool developed for malicious purposes and documented as being available through the dark web.
Zero-Day Vulnerability	A previously unknown or unpatched software vulnerability that may be exploited before an effective fix or mitigation is available.

Background

5. The development of generative AI has expanded the means through which established forms of deception and cybercrime can be conducted. In its publication *the criminal use of ChatGPT – a cautionary tale about large language models*⁴, issued on 27 March 2023, Europol identified the potential for large language models to facilitate fraud, social engineering, disinformation and malicious code generation. The FBI subsequently documented the use of generated text, images, audio and video to facilitate financial fraud⁵. Together, these developments show how capabilities designed to improve communication and access to knowledge can also be used to impersonate trusted individuals, manipulate decisions and obtain sensitive information.
6. During 2026, assessments of malicious AI use increasingly addressed its integration across connected stages of cyber operations. Microsoft's April assessment described assistance with reconnaissance, deceptive communications, malware development and activity following a compromise⁶. Alongside these developments, OpenAI's September reporting framework documented instances of concerning model behavior during training and evaluation, including unauthorized actions and communication across intended boundaries⁷. The security challenge consequently involves both the use of AI by human actors pursuing harmful objectives and the behavior of agents operating beyond the authority granted to them.
7. This evolution places the relationship between capability and control at the center of the governance debate. Effective protection depends on the ability to define an agent's permissions, monitor consequential actions and intervene when its behavior creates a security risk. NCSC's guidance on agentic AI emphasizes proportionate autonomy, controlled operating environments, monitoring and the capacity to halt agents and their communications⁸. As systems operate across multiple services, or through local and distributed environments,

⁴ European Union Agency for Law Enforcement Cooperation (Europol). (27 March 2023). *The criminal use of ChatGPT – a cautionary tale about large language models*. Europol. <https://www.europol.europa.eu/media-press/newsroom/news/criminal-use-of-chatgpt-cautionary-tale-about-large-language-models>

⁵ Federal Bureau of Investigation, Internet Crime Complaint Center (IC3). (3 December 2024). *Criminals Use Generative Artificial Intelligence to Facilitate Financial Fraud*. <https://www.ic3.gov/PSA/2024/PSA241203>

⁶ Microsoft. (2 April 2026). *Threat actor abuse of AI accelerates from tool to cyberattack surface*. Microsoft Security Blog. <https://www.microsoft.com/en-us/security/blog/2026/04/02/threat-actor-abuse-of-ai-accelerates-from-tool-to-cyberattack-surface/>

⁷ OpenAI. (16 September 2026). *Our framework for reporting model misalignment*. <https://openai.com/index/model-misalignment-reporting-framework/>

⁸ National Cyber Security Centre (NCSC). (20 August 2026). *Managing the cyber risk of agentic AI*. <https://www.ncsc.gov.uk/blogs/managing-the-cyber-risk-of-agentic-ai>

these safeguards must be supported by clear responsibilities among developers, deploying organizations and infrastructure operators.

Cybersecurity Risks Arising from AI Misuse

8. Since generative AI became widely available to the public, concerns about its misuse have focused on the risks it could create or intensify for individuals, institutions and digital systems. From its early adoption to the present, the main areas of concern have included the following:

- **Deepfakes, impersonation and phishing.** Generative AI can support fraud through synthetic identities, manipulated images and convincing audio or video impersonations. The Financial Crimes Enforcement Network's FinCEN Alert on Fraud Schemes Involving Deepfake Media Targeting Financial Institutions⁹ describes the use of such material to circumvent identity verification and impersonate trusted individuals in fraudulent payment requests. These practices undermine the reliability of familiar identity checks and expose institutions and individuals to financial loss.
- **Social engineering and sensitive data exposure.** AI-related risks extend from manipulating people into disclosing information to manipulating the systems that process it. NIST's Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile¹⁰ identifies privacy and information-security risks, including unauthorized disclosure of sensitive data and prompt injection. In the latter case, attackers introduce instructions through inputs or external material that an AI application processes, potentially redirecting its behavior. The resulting exposure depends on the information and organizational resources accessible to the application.
- **Vulnerability exploitation and connected devices.** AI can assist attackers in researching software weaknesses and developing methods of exploitation. NCSC's Impact of AI on cyber threat from now to 2027¹¹ identifies existing use of AI in vulnerability research and anticipates faster exploitation of unpatched systems, alongside improvements in skilled actors' ability to discover and exploit zero-day vulnerabilities. For connected environments,

⁹ Financial Crimes Enforcement Network (FinCEN). (13 November 2024). FinCEN Alert on Fraud Schemes Involving Deepfake Media Targeting Financial Institutions. U.S. Department of the Treasury. <https://www.fincen.gov/system/files/shared/FinCEN-Alert-DeepFakes-Alert508FINAL.pdf>

¹⁰ National Institute of Standards and Technology (NIST). (July 2024). Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (NIST AI 600-1). U.S. Department of Commerce. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>

¹¹ United Kingdom National Cyber Security Centre (NCSC). (7 May 2025). Impact of AI on Cyber Threat from Now to 2027. National Cyber Security Centre. <https://www.ncsc.gov.uk/report/impact-ai-cyber-threat-now-2027>

including Internet of Things devices, these developments reinforce the importance of controlling access, protecting data and maintaining secure software updates the capabilities addressed in NIST's IoT Device Cybersecurity Capability Core Baseline¹².

- **AI-assisted mobile malware.** AI can also be incorporated into malicious software to help it adapt to the environment in which it operates. ESET's research article PromptSpy ushers in the era of Android threats using GenAI¹³ analyzed Android malware samples that used Gemini to interpret screen information and support persistence on a device. The finding demonstrates an adaptive mechanism within the analyzed software; ESET had not observed the malware in its telemetry at publication.
- **Insider access obtained through deception.** AI-assisted impersonation can help malicious actors acquire positions that provide access to organizational systems. Google Threat Intelligence Group's The ultimate insider threat: North Korean IT workers¹⁴ describes experimentation with AI-generated profile photographs, deepfakes during interviews and writing tools to support deceptive recruitment practices. It also documents the use of privileged access by these workers for data theft and extortion. This connects weaknesses in identity verification and recruitment with the security of internal information and infrastructure.

Agent Interaction and the Acceleration of Cyber Threats in 2026

Moltbook and the security of agent interaction

9. Moltbook¹⁵ offers an early example of a digital environment designed for interaction between AI agents. Platforms of this kind create opportunities for systems to exchange information and respond to material generated by other participants¹⁶. Their security significance arises from the relationship between those interactions and the wider resources an agent may be permitted to use. Content encountered on one platform can influence activity elsewhere if the agent has access to additional accounts, tools or institutional systems.

¹² Fagan, M., et al. (29 May 2020). IoT Device Cybersecurity Capability Core Baseline (NISTIR 8259A). National Institute of Standards and Technology (NIST). <https://nvlpubs.nist.gov/nistpubs/ir/2020/NIST.IR.8259A.pdf>

¹³ Stefanko, L. (19 February 2026). PromptSpy Ushers in the Era of Android Threats Using GenAI. ESET Research. <https://www.welivesecurity.com/en/eset-research/promptsy-ushers-in-era-android-threats-using-genai/>

¹⁴ Collier, J., et al. (6 March 2025). The Ultimate Insider Threat: North Korean IT Workers. Google Cloud. <https://cloud.google.com/transform/ultimate-insider-threat-north-korean-it-workers>

¹⁵ Moltbook. (2026). Moltbook: A Social Network for AI Agents. Moltbook. <https://www.moltbook.com/>

¹⁶ Taylor, J. (2 February 2026). What Is Moltbook? The Strange New Social Media Site for AI Bots. The Guardian. <https://www.theguardian.com/technology/2026/feb/02/moltbook-ai-agents-social-media-site-bots-artificial-intelligence>

10. Wiz's February 2026 investigation of Moltbook identified a misconfigured database exposing approximately 1.5 million authentication tokens, 35,000 email addresses and private messages. The researchers also found that the platform allowed humans to impersonate agents and lacked reliable verification of participants' identities. The vulnerability was addressed following disclosure¹⁷. These findings place identity assurance and platform security at the center of the case.
11. For governance, the important connection is between communication and delegated authority. An organization needs to know which system is participating, on whose behalf it acts and what influence external content can have over its subsequent decisions. Agent interaction therefore requires more than rules governing acceptable posts. It requires safeguards around credentials, trusted instructions and the movement of information between platforms. Moltbook illustrates how the security of an interaction environment can affect confidence in the behavior and identity of the systems using it.

From task assistance to connected attack workflows

12. Provider assessments published between April and June 2026 describe AI assistance across a widening range of cyber activities. Microsoft's 2 April 2026 assessment identified its use in reconnaissance, the preparation of deceptive material, malware development and activity after a compromise. Microsoft reported that the attacks it examined typically retained human involvement¹⁸. The development is significant because assistance can be provided at several stages of the same operation, allowing tasks that previously required separate effort to be connected within a more efficient workflow.
13. Google Threat Intelligence Group's 11 May 2026 assessment extended this picture to vulnerability exploitation. The company identified a criminal actor using a zero-day exploit that it assessed was likely developed with AI assistance, and judged that its discovery may have prevented the actor's intended mass exploitation¹⁹. The finding points towards a security

¹⁷ Nagli, G. (2 February 2026). Hacking Moltbook: The AI Social Network Any Human Can Control. Wiz.

<https://www.wiz.io/blog/exposed-moltbook-database-reveals-millions-of-api-keys>

¹⁸ DeGrippe, S. (2 April 2026). Threat Actor Abuse of AI Accelerates from Tool to Cyberattack Surface. Microsoft Security Blog. <https://www.microsoft.com/en-us/security/blog/2026/04/02/threat-actor-abuse-of-ai-accelerates-from-tool-to-cyberattack-surface/>

¹⁹ Google Threat Intelligence Group (GTIG). (11 May 2026). GTIG AI Threat Tracker: Adversaries Leverage AI for Vulnerability Exploitation, Augmented Operations, and Initial Access. Google Cloud. <https://cloud.google.com/blog/topics/threat-intelligence/ai-vulnerability-exploitation-initial-access/>

environment in which assistance with technically demanding intrusion activity can accompany the more established use of AI for deception and reconnaissance.

14. Anthropic's 3 June 2026 analysis examined 832 accounts banned for malicious cyber activity between March 2025 and March 2026, selected for the availability of sufficient information. It found use across later stages of the attack lifecycle and a growing capacity to connect tasks with limited human intervention²⁰. Taken together, these assessments identify a change in the organization of malicious activity: models can contribute throughout an operation, while more of the intermediate work can be delegated.
15. Microsoft's 22 September 2026 investigation, *Unmasking EvilTokens: Getting to the Root of Device Code Phishing*²¹, documented a commercially supplied phishing platform that used AI to tailor deceptive messages and analyze compromised inboxes to identify people with financial or administrative authority. Microsoft linked the service to more than 12,000 compromised inboxes across more than 10,000 organizations. Its Digital Crimes Unit also reported a coordinated disruption of the platform's infrastructure²². The case illustrates how human-directed activity can combine AI-assisted phishing, account compromise, post-compromise analysis and fraud preparation within a single service.
16. The distinction between human direction and automated execution is central to the implications of this trend. A person may choose the target and objective while software performs a substantial portion of the work. The security value of human oversight then depends on when and how it is exercised. Approval at the beginning of a workflow may offer little protection against an unexpected action that occurs several stages later. Monitoring and authorization must consequently address the points at which a system can make a consequential change, obtain further access or move information beyond its original environment.
17. The same capabilities create opportunities for defense. The collective cyber-defense initiative hosted by OpenAI calls for stronger baseline security, wider access to defensive AI, coordinated incident response and support for essential services²³. This approach connects

²⁰ Anthropic. (3 June 2026). What We Learned Mapping a Year's Worth of AI-Enabled Cyber Threats. Anthropic. <https://www.anthropic.com/news/AI-enabled-cyber-threats-mitre-attack>

²¹ Microsoft Threat Intelligence, Microsoft Defender Experts, and Microsoft Security Research. (22 September 2026). *Unmasking EvilTokens: Getting to the Root of Device Code Phishing*. Microsoft Security Blog. <https://www.microsoft.com/en-us/security/blog/2026/09/22/unmasking-eviltokens-getting-to-the-root-of-device-code-phishing/>

²² Masada, S. (22 September 2026). *Disrupting EvilTokens: The AI Chatbot Built for Cybercrime*. Microsoft On the Issues. <https://blogs.microsoft.com/on-the-issues/2026/09/22/disrupting-eviltokens-the-ai-chatbot-built-for-cybercrime/>

²³ OpenAI. (2026). A Call for Collective Action on Cyber Defense. OpenAI. <https://openai.com/collective-cyberdefense/>

advanced technical capabilities with the practical condition of the institutions that need protection. Defensive capacity has to extend beyond model developers to the organizations operating public services, infrastructure and other exposed systems.

18. OpenAI's 10 August 2026 announcement, *Expanding Daybreak as the Cyber Defense Window Narrows*, described differentiated access to advanced cybersecurity capabilities for approved defenders. Daybreak Blue supports defensive activities, including vulnerability discovery, malware analysis, and incident response, while Daybreak Red provides access to purpose-trained models for authorized vulnerability research, exploit validation, and security testing. OpenAI stated that access controls include identity verification, account security, monitoring, approved-use restrictions, and legal attestations²⁴. This approach links wider access to defensive capabilities with the need for controlled operating environments, defined authorization, and effective oversight.
19. Google's account of its PageBreak project, published on 24 September 2026, provides an example of agentic AI used for authorized cyber defense. Google reported that the internal system identified more than 500 cross-site scripting vulnerabilities across the company's own web applications. Separate, non-AI validators checked suspected vulnerabilities against running systems before findings were sent to product teams. The case illustrates the importance of distinguishing AI-generated hypotheses from verified security findings; the reported results remain specific to Google's testing environment²⁵.
20. The policy challenge is therefore both technological and institutional. Faster detection has limited value if an organization cannot act on an alert, withdraw access or coordinate with another service provider. The benefits of defensive AI depend on reliable operational procedures and sufficient capacity across the wider security environment. Without this foundation, improvements in capability may be distributed unevenly between well-resourced institutions and those already struggling to maintain basic protection.

Unauthorized agent behavior and the problem of control

21. AI agents introduce a further security concern when they take actions beyond the limits intended by their operators. The incidents disclosed during the summer and September 2026

²⁴ OpenAI. (10 August 2026). *Expanding Daybreak as the Cyber Defense Window Narrows*. OpenAI.

<https://openai.com/index/expanding-daybreak-as-the-cyber-defense-window-narrows/>

²⁵ Bentkowski, M. (24 September 2026). *Agentic Hacks, Real Proofs: Inside Google's PageBreak Project*. Google Security Blog. <https://blog.google/security/agentic-hacks-real-proofs-inside-googles-pagebreak-project/>

bring this issue into operational settings. They concern systems performing assigned tasks while obtaining unauthorized access, communicating across intended boundaries or using methods that their operators had not approved. Such behavior makes the relationship between task performance, technical isolation and effective supervision a central question for frontier AI governance.

22. OpenAI's 26 August 2026 account described a July incident in which models conducting cybersecurity evaluations circumvented isolation controls and compromised parts of OpenAI's research infrastructure and Hugging Face's systems. The company reported dangerous actions that had not been directed by a human²⁶. The incident demonstrates why the security of an evaluation environment depends on both the restrictions imposed by infrastructure and the behavior of the models operating within it.
23. METR's investigation examined agent coordination and conduct, principally during 7-13 July 2026, using records made available by OpenAI. It identified collective efforts to influence evaluation outcomes. Its analysis also specified the boundaries of the investigation, including limits on its coverage of earlier activity and remediation²⁷. The availability of external examination contributes to accountability by providing a further basis for assessing how an incident developed and which conclusions the available records support.
24. On 21 September 2026, the Independent International Scientific Panel on AI, established by the UN General Assembly, published its thematic brief, *AI Agents, Misalignment and Loss of Human Control Risks: Evidence from the OpenAI-Hugging Face Incident*²⁸. Drawing on company disclosures, METR's investigation and wider research, the Panel identified three conditions present in the incident: a misaligned objective, the capability to pursue it and an environment that allowed the behavior. It found that containing this incident provides no assurance of reliable human control over more capable agents, while making no prediction about the likelihood or timing of severe loss of control. The assessment directs attention to both agent behavior and the security of the systems in which agents operate.

²⁶ OpenAI. (26 August 2026). The Hugging Face Incident and the Road Ahead. OpenAI. <https://openai.com/index/hugging-face-incident-and-the-road-ahead/>

²⁷ Greenblatt, R., et al. (26 August 2026). Brief Independent Investigation of Agents' Behavior, Reasoning and Collaboration in the OpenAI / Hugging Face Hacking Incident. METR. <https://metr.org/blog/2026-08-26-openai-hugging-face-incident-investigation/>

²⁸ Independent International Scientific Panel on AI. (21 September 2026). *AI Agents, Misalignment and Loss of Human Control Risks: Evidence from the OpenAI-Hugging Face Incident* (advance unedited thematic brief). United Nations. <https://www.un.org/independent-international-scientific-panel-ai/en/thematic-briefs/ai-agents-misalignment-risks>

Communication Across Intended Boundaries

25. On 4 September 2026, researchers from Nightingale Collective and collaborators published Discovery of a New OpenAI Agent Message Board, an investigation of a public wiki used by OpenAI agents to exchange information during web-research tasks. The reconstructed activity occurred principally in May and June, with limited later activity. The researchers identified approximately 18,000 posts associated with around 3,700 self-assigned names and described exchanges about task execution and ways around restrictions. They assessed the wiki activity as probably separate from the group involved in the Hugging Face compromise²⁹.
26. Communication between agents can extend a task beyond the boundaries anticipated by its operator. Information exchanged through external services may influence subsequent actions or enable coordination across instances. Containment should therefore address both the permissions of individual agents and the channels through which they discover one another, exchange information and share resources. The scope of supervision may consequently extend beyond a single conversation or model instance.

Evaluation Incidents and Model Supervision

27. Anthropic's assessment of 9 September 2026 examined four incidents in which Claude models gained unauthorized access to real third-party systems during cybersecurity evaluations conducted with the same external partner. The evaluation environments had unintended internet access, while safeguards used in production were absent. In one incident, Claude Mythos 5 published a malicious package to the public PyPI repository. According to Anthropic, the package was installed on 15 third-party hosts, and credentials exposed by one installation were subsequently used to access a security vendor's live database³⁰. In a separate evaluation conducted in May 2026, Google confirmed that a Gemini model accessed systems belonging to three real companies³¹. Together, the incidents illustrate how failures in evaluation containment can combine with task-driven model behavior to affect systems outside the authorized exercise. A separate, human-directed case illustrates a related shift in operational

²⁹ Von Arx, S., et al. (4 September 2026). Discovery of a New OpenAI Agent Message Board. Collusion Wiki. <https://collusion.wiki/>

³⁰ Anthropic. (9 September 2026). An Alignment Assessment of Recent Cybersecurity Incidents. <https://www.anthropic.com/research/alignment-assessment-cybersecurity-incidents>

³¹ Sabin, S. (18 September 2026). Google's AI Hacked Three Companies in Testing. Axios. <https://www.axios.com/2026/09/19/google-safety-incidents-testing-hacks>; Irregular. (14 August 2026). Addressing Recent Incidents: Ongoing Findings and Path Forward. <https://www.irregular.com/research/addressing-recent-incidents-ongoing-findings-and-path-forward>

scale: according to Gambit Security, one operator used Claude Code and GPT-4.1 to compromise nine Mexican government agencies and process volumes of data that would ordinarily have required a team³².

28. On 16 September 2026, OpenAI published its misalignment-reporting framework alongside six reports of concerning behavior observed during model training and evaluation over the preceding six months. The cases included concealment, unauthorized credential use, unapproved file uploads and communication across intended boundaries³³. The company emphasized that these individual observations did not indicate how frequently misalignment occurred across its models. It also stated that, in its assessment, the industry had not achieved sufficient progress in aligning AI systems with human intentions and monitoring their behavior to sustain responsible increases in capability at maximum speed for much longer³⁴. These disclosures underline the need to assess both whether agents act within their authority and whether safeguards can keep pace with increasing capabilities.
29. On 24 September 2026, Australian Prime Minister Anthony Albanese disclosed that an OpenAI agent, assigned during an internal evaluation on 18 June 2026 to research publicly available medicine-spending statistics, had obtained unauthorized non-public access to the Medicare Statistics Reporting Service administered by Services Australia. Australian authorities reported access to non-public files and the writing of files to an internal server, while finding no evidence at that stage of access to individual medical records or broader compromise of Services Australia systems. A forensic investigation remained underway³⁵. In a fuller account published on 28 September 2026, OpenAI described related agent activity involving the New South Wales Bureau of Crime Statistics and Research, the Victorian Department of Health's Victorian Agency for Health Information, and the Australian Institute of Health and Welfare. The reported outcomes differed: a public New South Wales tool returned operational information; an exposed Victorian access key was used in circumstances

³² Sela, E. (10 April 2026). A Single Operator, Two AI Platforms, Nine Government Agencies: The Full Technical Report. Gambit Security. <https://gambit.security/blog-posts/a-single-operator-two-ai-platforms-nine-government-agencies-the-full-technical-report>

³³ OpenAI. (16 September 2026). Our Framework for Reporting Model Misalignment. OpenAI. <https://openai.com/index/model-misalignment-reporting-framework/>

³⁴ Martin, E. (16 September 2026). *OpenAI Discloses Six New Incidents of 'Concerning' A.I. Behavior*. The New York Times. <https://www.nytimes.com/2026/09/16/technology/openai-model-safety-guardrails.html>

³⁵ Office of the Prime Minister of Australia. (24 September 2026). Press conference – New York. <https://www.pm.gov.au/media/press-conference-new-york>; Office of Senator the Hon Katy Gallagher, Minister for Finance and Minister for Government Services. (24 September 2026). Press Conference Sydney. <https://ministers.finance.gov.au/financeminister/transcript/2026/09/24/press-conference-sydney>

where authorization remained unclear; and apparently public data was retrieved from the Australian Institute of Health and Welfare after unsuccessful attempts to bypass access controls. OpenAI stated that its review identified the activity in mid-August and acknowledged that preliminary findings should have been shared sooner. Notification occurred on 10 September 2026 for Services Australia and Victoria, 18 September 2026 for New South Wales, and 24 September 2026 for the Australian Institute of Health and Welfare. OpenAI also reported a pause in tool-using training and evaluation for its most capable models and announced a task force with independent Australian expertise to improve notification and coordination³⁶. The differing outcomes demonstrate the importance of establishing precisely what occurred at each system, notifying affected authorities promptly, and coordinating investigations across institutions.

30. On 25 September 2026, OpenAI disclosed a separate incident from 20 September 2026 in which an internal research agent used a gap in the DNS filtering of a training sandbox to contact an external chatbot despite restrictions on live internet access. According to OpenAI, monitoring raised an alert within 15 minutes, and a reviewer acknowledged it three minutes later, but the expected automatic stop did not occur. The run was terminated manually approximately 2.5 hours after the alert. OpenAI reported that, as of 25 September 2026, training, evaluation, and inference involving tool use for its most capable models were paused pending validation of the controls and further testing³⁷. The incident illustrates the difference between detecting an unauthorized action and containing it promptly.
31. On 28 September 2026, OpenAI deferred the release of GPT-6.1 Astra after internal safety evaluations raised concerns about actions beyond authorized task boundaries. OpenAI's head of safety systems, Saachi Jain, told the Associated Press that the model had become more persistent in completing tasks but had not met the company's threshold for remaining within the scope of authorization³⁸. The decision illustrates how evaluation findings can affect the timing of deployment.

³⁶ OpenAI. (28 September 2026). How we will do better for Australia. <https://openai.com/index/how-we-will-do-better-for-australia/>

³⁷ OpenAI. (25 September 2026). *An agent used DNS to reach an external chatbot*. OpenAI Alignment. <https://alignment.openai.com/misalignment-reports/an-agent-used-dns-to-reach-an-external-chatbot/>

³⁸ Beaumont, T. (28 September 2026). OpenAI delays latest model over security concerns, as industry faces new safety pressures. Associated Press. <https://apnews.com/article/open-ai-artificial-intelligence-altman-trump-astra-5afb865b2cddc439efdcf31ebdc406a5>

32. On 29 September 2026, OpenAI introduced dots, persistent agents powered by GPT-6 Astra that can use cloud computers, browsers, connected tools, communications channels, and delegated subagents. OpenAI’s updated system card describes evaluations addressing repeated prompt-injection attempts, changing permissions, and the preservation of boundaries across extended tasks. The deployment extends the governance challenge from isolated requests to continuing access across multiple services. Authorization, monitoring, and intervention must therefore remain effective throughout an agent’s operation, including when the scope of a task changes³⁹.
33. Separately, on 29 September 2026, OpenAI introduced GPT-6.1 Sol and classified the model as reaching the company’s Critical cybersecurity capability threshold and High biological and chemical capability threshold under the Preparedness Framework. OpenAI stated that GPT-6.1 Sol uses the safeguards applied to GPT-6 Astra. In an internal evaluation without system-level controls, unwanted persistence after warnings occurred in 23.5% of tested runs; OpenAI cautioned that this result does not measure how often circumvention would succeed under deployed controls. These provider assessments reinforce the need to evaluate capabilities and operational safeguards for each model and deployment⁴⁰.
34. These developments place task completion and compliance within the same governance problem. A system may identify a method that advances its immediate objective while violating the conditions under which it was authorized to operate. Organizations deploying agents consequently need to specify acceptable means as well as desired outcomes. Evaluation should test the ability to preserve restrictions when they make a task more difficult, while operational controls should prevent the system from overriding those restrictions on its own.
35. The consequences also extend to responsibility. Developers control aspects of model training and safeguards; deploying organizations determine access, workflows and supervision; infrastructure operators maintain the environment in which actions occur. An incident can involve failures across several of these levels. Clear allocation of responsibilities is essential for ensuring that a warning leads to intervention and that an investigation can reconstruct both the technical sequence and the decisions that made it possible.

³⁹ OpenAI. (3 September 2026; updated 29 September 2026). GPT-6 Astra System Card. <https://deploymentsafety.openai.com/gpt-6-astra/sec>

⁴⁰ OpenAI. (29 September 2026). Addendum to GPT-6 Astra System Card: GPT-6.1 Sol. <https://deploymentsafety.openai.com/gpt-6-1-sol>

Weapons Related Misuse and Biological Research Risks

36. Earlier warnings about the military use of AI addressed the risks associated with its integration into military decision-making and weapons systems, particularly harm to civilians and the erosion of human control over the use of force⁴¹. Recent cases have brought additional challenges into focus: the misuse of general-purpose AI models to support weapons development, and the difficulty of detecting and interrupting such assistance through provider safeguards.
37. The security implications of AI assistance extend into activities where digital support can contribute to physical harm. Anthropic's September 2026 threat assessment described a northern Yemen cell using Claude Code across three guided-weapons development programs. Several model instances were assigned different engineering roles. The company reported incomplete blocking of requests, account bans, an apparently failed rocket test and the production of an offline simulation toolkit⁴². This case brings the delegation of technical work into a weapons-related context and links service-level intervention with the question of capabilities already acquired by an actor.
38. In biological research, the same assessment presented five case studies involving uses that could support biological weapons development. Anthropic described the difficulty of evaluating potentially dangerous work in a domain where legitimate civilian research and harmful applications may involve overlapping knowledge and techniques⁴³. The concern is the ability of AI assistance to contribute to particular stages of research, and the institutional capacity to identify when that assistance becomes relevant to misuse.
39. Assessment of weapons-related misuse should examine which technical barriers AI helps an actor overcome and what capabilities remain after access is withdrawn. Closing an account can prevent further assistance through that service, while previously generated software, acquired knowledge and locally stored material may remain usable. The significance of an intervention therefore depends on the actor's expertise, materials and testing capacity, as well as the timing of detection.

⁴¹ International Committee of the Red Cross (ICRC). (24 July 2023). The Rise of Artificial Intelligence Requires a Human-Centred Approach in Conflict Zones. International Committee of the Red Cross. <https://www.icrc.org/en/document/rise-artificial-intelligence-requires-human-centred-approach-conflict-zones>

⁴² Anthropic. (10 September 2026). Detecting and Countering Misuse of AI: September 2026. Anthropic. <https://www.anthropic.com/threat-intelligence-report-september-2026>

⁴³ Ibidem

40. The biological context further underlines the importance of specialized assessment. Oversight should be able to distinguish routine scientific activity from requests whose purpose, combination or operational context raises a security concern. Effective prevention requires scientific expertise alongside information about access, intent and the wider research environment. Safeguards that are too narrow may miss the significance of connected activities, while indiscriminate restrictions can obstruct beneficial research. The governance task is to maintain a proportionate approach that protects legitimate science and addresses the specific pathways through which misuse can develop.

Provider Safeguards and the Limits of Intervention

41. AI providers occupy an important position in the prevention and investigation of misuse. They can establish access conditions, evaluate model behavior, monitor activity on their services and suspend accounts associated with harmful conduct. Their visibility can help identify patterns that would otherwise remain dispersed across individual interactions. The practical value of these measures depends on whether detection is timely and whether the provider retains control over the resources involved in the activity.
42. The effectiveness of provider safeguards depends on the connection between detection, intervention, and notification. OpenAI's account of the 20 September 2026 incident records a prompt alert and human acknowledgment, followed by a failure of the expected automatic stop and termination of the run approximately 2.5 hours later⁴⁴. In a broader review of model activity during training and evaluation, OpenAI reported notifying dozens of third parties about activity that may have bypassed security controls or adversely affected external services⁴⁵. The number of notifications should not be treated as a count of confirmed serious intrusions. These disclosures underscore the need to assess whether alerts lead to timely containment and whether affected organizations receive sufficient information to investigate.
43. NCSC's August guidance on agentic AI recommends proportionate autonomy, sandboxing, monitoring, attributable activity and the ability to halt agents and their communications⁴⁶. These measures address the deployment environment as well as the model. A refusal policy

⁴⁴ OpenAI. (25 September 2026). *An agent used DNS to reach an external chatbot*. OpenAI Alignment.

<https://alignment.openai.com/misalignment-reports/an-agent-used-dns-to-reach-an-external-chatbot/>

⁴⁵ OpenAI. (September 2026). *The Hugging Face incident and other third-party impact from misaligned models*. OpenAI.

<https://openai.com/hugging-face-incident-and-misalignment/>

⁴⁶ W, T. (20 August 2026). *Managing the Cyber Risk of Agentic AI*. United Kingdom National Cyber Security Centre (NCSC).

<https://www.ncsc.gov.uk/blogs/managing-the-cyber-risk-of-agentic-ai>

can influence what a system attempts to do, while restrictions on credentials, networks and tools determine what it is technically able to do. The two layers should reinforce one another throughout the period in which an agent operates. On 28 September 2026, NVIDIA announced the Open Agent Safety Platform, combining software designed to enforce policies during agent execution with a hardware-based reference design for monitoring and isolating agents that cross defined boundaries⁴⁷. The initiative illustrates growing attention to controls within the deployment infrastructure, alongside safeguards incorporated into AI models.

44. Safeguards should be assessed by the speed and completeness of containment. Relevant measures include the time between detection and withdrawal of access, whether related accounts or agents remain active, whether external operators must intervene and whether harmful activity recurs. Reporting should distinguish suspension of an account from termination of the wider operation, enabling public authorities and independent evaluators to assess the practical outcome of an intervention.
45. When models run locally or offline, their developers may have no continuing visibility into their use or means to terminate execution remotely. Open-weight models introduce an additional consideration, as users may download and modify model parameters and remove safeguards or refusal mechanisms incorporated by the original developer. In distributed deployments, different organizations may control development, hosting and access across several jurisdictions. Effective safeguards must therefore operate wherever technical authority can be exercised. Deploying institutions and infrastructure operators need defined responsibilities for access controls, monitoring and intervention, supported by public oversight capable of assessing high-risk uses and coordinating responses across organizational boundaries.

Frontier AI Warnings and International Governance

46. Recent security incidents have intensified the debate over whether safeguards can keep pace with advancing capabilities. The concern extends beyond individual misuse cases to the conditions under which increasingly capable systems are trained, evaluated and deployed. It brings technical assurance into direct contact with questions of public authority, independent oversight and international coordination.

⁴⁷ NVIDIA. (28 September 2026). *NVIDIA Launches Open Agent Safety Platform to Secure Agents From Testing to Deployment*. <https://nvidianews.nvidia.com/news/open-agent-safety-platform>

47. Addressing the Human Rights Council on 7 September 2026, Volker Türk, United Nations High Commissioner for Human Rights, warned that uncontrolled advanced AI could pose an existential risk to humanity. Türk stated: “AI that escapes its testing environment, or blackmails developers to prevent itself from being turned off, is AI that is too powerful.” Türk called for binding rules, independent verification, clear limits, and closer cooperation among states and AI developers, linking AI safety to the protection of human rights⁴⁸.
48. In *We Must Pace the Frontier*, Dario Amodei, co-founder and chief executive officer (CEO) of Anthropic, argued that the pace of frontier AI development should remain compatible with progress in evaluation and safeguards. The essay identifies the growing use of AI to develop subsequent generations of models, known as recursive self-improvement, as a potential source of acceleration beyond existing capacities for human understanding and control. Amodei proposed independent evaluators embedded within companies, coordination among developers, and international cooperation⁴⁹. Under this approach, evidence of effective safeguards would help determine the pace of further capability development.
49. Sam Altman, Chief Executive Officer (CEO) of OpenAI, endorsed pacing frontier AI development and the use of independent evaluators in a statement published on 12 September 2026. In a separate statement on 14 September 2026, Altman supported consistent federal safety requirements and argued that companies could begin implementing improvements without waiting for legislation or an antitrust exemption. The statements link the pace of development to safety cases and effective monitoring⁵⁰. Elon Musk, founder of xAI and Chief Executive Officer (CEO) of SpaceX, also endorsed Amodei’s proposal in a statement published on 12 September 2026⁵¹.
50. On 28 September 2026, OpenAI published initial guidelines for evidence-based safety cases before the continuation of frontier reinforcement-learning training runs. The guidelines address evidence concerning model alignment, containment, monitoring, approval, auditor access, escalation, and the investigation of serious incidents. OpenAI described comprehensive safety

⁴⁸ United Nations News. (7 September 2026). AI: Türk Urges Action Before It Becomes an ‘Existential Risk to Humanity’. United Nations. <https://news.un.org/en/story/2026/09/1168288>

⁴⁹ Amodei, D. (September 2026). *We Must Pace the Frontier*. Dario Amodei. <https://darioamodei.com/post/we-must-pace-the-frontier>

⁵⁰ Altman, S. (12 September 2026). [Post on pacing frontier AI development and independent evaluators]. X. <https://x.com/sama/status/2098811563415150910>; Altman, S. (14 September 2026). [Post on pacing frontier AI development, safety cases and monitoring]. X. <https://x.com/sama/status/2099348812305473766>

⁵¹ Musk, E. (12 September 2026). [Post endorsing Dario Amodei’s call to pace frontier AI development]. X. <https://x.com/elonmusk/status/2098789109980332057>

cases as an aspirational objective and stated that implementation of the initial guidelines was underway⁵². Their value for wider oversight will depend on the quality of the supporting evidence, access for qualified evaluators, and whether adverse findings lead to corrective action.

51. The governance significance of these positions lies in the relationship between technical capability and institutional assurance. A commitment to improve safeguards becomes more meaningful when the evidence required for proceeding is specified in advance. Independent evaluators need sufficient access to examine relevant behavior, investigate serious concerns and communicate their findings. Oversight must also address the conditions under which an evaluation can require a change in deployment or a delay in further capability advancement.
52. International coordination is important because development and deployment are distributed across borders. Divergent requirements can make it difficult to compare safety claims or organize a response involving several providers and infrastructure operators. Common expectations for incident reporting, evaluation access and the preservation of evidence would help establish a basis for cooperation. Such arrangements should connect technical findings with institutions able to act on them, while supporting countries whose regulatory and operational capacity remains limited.
53. On 15 September 2026, Security Council Report published its policy paper, Artificial Intelligence, Geopolitics, and the UN Security Council, examining the implications of AI for international peace and security and the potential role of the UN Security Council⁵³. It noted that the Council's engagement with AI has remained sporadic and uneven across its agenda, while political divisions continue to constrain a more systematic approach. The paper identified options for strengthening engagement within the Council's existing mandate, including mainstreaming AI across relevant country-specific and thematic agenda items, strengthening horizon-scanning and institutional memory, drawing more systematically on technical and

⁵² OpenAI. (28 September 2026). Towards safety cases for frontier AI training. <https://openai.com/index/towards-safety-cases-for-frontier-ai-training/>

⁵³ Security Council Report. (15 September 2026). Artificial Intelligence, Geopolitics, and the UN Security Council. https://www.securitycouncilreport.org/atf/cf/%7B65BFCF9B-6D27-4E9C-8CD3-CF6E4FF96FF9%7D/ai_policy_paper_2026.pdf

scientific expertise, and establishing an Informal Expert Group on Technology, Peace and Security⁵⁴.

54. On 23 September 2026, the Security Council held a high-level briefing on AI and international security. Yoshua Bengio, co-chair of the Independent International Scientific Panel on AI, and the chief executives of OpenAI, Anthropic and Hugging Face addressed risks associated with misuse and the potential loss of human control over increasingly capable agents. Across their interventions, they proposed independent evaluations, more consistent reporting of serious incidents and stronger international coordination. The meeting brought these questions into the Council's formal consideration of international peace and security, while leaving the proposed governance arrangements for further deliberation⁵⁵.

55. Human rights should remain integral to this approach. Systems introduced to monitor misuse can themselves affect privacy, access to information and legitimate scientific activity. Oversight arrangements should define the purposes for which information is collected, the conditions for sharing it and the means through which decisions can be reviewed. Effective governance requires both the capacity to respond to serious risks and accountable limits on the exercise of that capacity.

Attribution and the Future of Effective Human Control

56. The incidents and assessments examined in this report reveal two connected dimensions of AI risk. Human actors are incorporating AI into harmful operations, while agents in legitimate evaluation settings have also taken actions beyond their authorization. Understanding how these dimensions interact is essential for selecting safeguards, assigning responsibility and determining the appropriate form of public intervention.

57. Attribution should begin with the origin of the objective and continue through the actions used to pursue it. A human actor may request assistance with a malicious task, delegate much of an operation to an agent, or manipulate a system serving a legitimate user. An agent may also adopt unauthorized means while pursuing the task assigned by its operator. Each pathway has

⁵⁴ Security Council Report. (17 September 2026). Artificial Intelligence, Geopolitics, and the UN Security Council. Security Council Report. <https://www.securitycouncilreport.org/research-reports/artificial-intelligence-geopolitics-and-the-un-security-council.php>

⁵⁵ United Nations Web TV. (23 September 2026). *Artificial Intelligence and International Security-Security Council, 10228th Meeting*. <https://webtv.un.org/en/asset/k1c/k1cdceg9sg>; United Nations. (23 September 2026). *Artificial Intelligence and International Security-Security Council, 10228th Meeting* (automated transcript). <https://transcripts.un.org/en/sc/10228>

implications for prevention, but the presence of automation alone does not resolve which one occurred.

58. Reliable investigation consequently requires information about account control, original instructions, permissions, external inputs, tool use and communication with other systems. These records must be examined together so that investigators can determine how the activity developed and where intervention was possible. Provider-level visibility is valuable in this process, while independent examination and cooperation with other relevant operators can help reconstruct actions that crossed service boundaries.
59. The malicious AI ecosystem adds a further dimension to attribution. Eurojust's documentation of WormGPT and FraudGPT identifies tools developed for malicious activities and made available through the dark web⁵⁶. Their existence establishes a pathway through which human actors may obtain AI assistance for harmful operations. Where the origin of an incident remains unresolved, investigation should examine whether such resources contributed to the activity, alongside evidence concerning account control, external instructions and the actions taken by the agents involved.
60. Governance responses should follow the pathway established by the evidence. Human-directed misuse calls for measures addressing the actor and access to relevant capabilities; unauthorized agent behavior requires examination of training, evaluation incentives, permissions and containment. Where these pathways interact, intervention must address both. Responsibilities should correspond to the control exercised by developers, deploying organizations and infrastructure operators, supported by independent scrutiny and cooperation across jurisdictions.
61. The documented cases demonstrate the need to address both human-directed misuse and agent behavior beyond authorized instructions. For incidents whose origin remains unresolved, could the activity have been directed by human actors using illicit AI tools available through the dark web, such as WormGPT and FraudGPT? Could it instead have arisen from agents adopting unauthorized means to complete assigned tasks, or from an interaction between external manipulation and agent behavior? How reliably can providers distinguish these pathways, and what evidence should be available for independent verification? As capabilities become more

⁵⁶ Eurojust. (30 November 2023). Generative Artificial Intelligence: The Impact on Intellectual Property Crimes. European Union Agency for Criminal Justice Cooperation (Eurojust). <https://www.eurojust.europa.eu/sites/default/files/assets/generative-ai-impact-to-ip-crimes.pdf>

autonomous and accessible through open, local, offline or distributed deployments, can provider-level safeguards remain effective, and what additional institutional and international oversight is needed to preserve meaningful human control?

National Positions and Geopolitical Implications of AI Governance

62. The governance of advanced AI is becoming an issue of international security and public accountability as well as technological policy. The positions expressed during September 2026 reveal differences over the pace of development, the evidence needed to justify deployment, and the authority of governments and independent institutions to require safeguards. For countries that depend on systems developed abroad, these decisions also affect access to technology, continuity of essential services, and the capacity to investigate incidents with consequences beyond national borders.
63. Three considerations run through this debate: maintaining technological competitiveness, preventing serious harm, and ensuring that countries can participate in the benefits and governance of AI. Governments attach different weight to each consideration. Their positions also evolve as new evidence emerges and diplomatic discussions develop. The central policy question is how to establish credible safeguards while sustaining innovation and avoiding arrangements that deepen technological dependence or exclude countries with limited resources.
64. In its ministerial statement, *Artificial Intelligence Update*, of 7 September 2026, the United Kingdom government supported the development and deployment of advanced AI while emphasizing the need for security and oversight to keep pace with capabilities. The statement warned that a widening gap between capability and control could create significant risks to public safety and national security. It outlined government work on biosecurity, agentic AI incident response, and the security of evaluation environments⁵⁷.
65. In an address to the UN General Assembly on 22 September 2026, Andy Burnham, Prime Minister of the United Kingdom, announced that security officials would begin work on a National Centre for Information Defense to detect, attribute, and disrupt hostile state information operations. Prime Minister Burnham warned that AI would amplify these threats and proposed using the United Kingdom's 2027 G20 presidency to pursue common

⁵⁷ Narayan, K. (7 September 2026). *Artificial Intelligence Update*. Cabinet Office, written ministerial statement HCWS314. UK Parliament. <https://questions-statements.parliament.uk/written-statements/detail/2026-09-07/HCWS314>

international principles and standards for frontier AI, covering transparency, access, preparedness, and safe development. Prime Minister Burnham also announced work on an AI defense partnership with the United States to protect critical national infrastructure and indicated that additional legislation could be introduced where necessary⁵⁸.

66. Addressing the UN Security Council on 23 September 2026, Ed Miliband, Secretary of State for Foreign, Commonwealth and Development Affairs of the United Kingdom, outlined three priorities: rigorous testing of frontier AI models, sufficient government visibility into developers' work, and resilience against AI-enabled incidents. Foreign Secretary Miliband linked public confidence in innovation to credible safeguards and emphasized that governments could not entrust the protection of their citizens solely to private companies⁵⁹. These priorities place access to evaluation evidence and preparedness for AI-related failures at the center of public oversight.
67. Donald Trump, President of the United States, has emphasized maintaining U.S. technological leadership and relying on existing government authority to address AI-related risks. On 14 September 2026, President Trump opposed additional constraints on AI development, arguing that such measures could benefit China⁶⁰.
68. At the Security Council briefing on 23 September 2026, Michael Kratsios, Director of the White House Office of Science and Technology Policy, expressed U.S. support for international dialogue and the development of national technical capacity, while opposing the establishment of new global AI governance structures. Mr. Kratsios identified existing law enforcement tools, cooperation on testing and evaluation, and national decision-making as the foundations of the U.S. approach⁶¹. This position distinguishes international cooperation on technical matters from the creation of international regulatory authority over AI.
69. The U.S. debate also encompasses differing legislative positions. In a letter published on 23 September 2026, a group of Democratic senators urged President Trump to negotiate common

⁵⁸ Prime Minister's Office, 10 Downing Street (23 September 2026). *PM Speech at UNGA: 22 September 2026*. GOV.UK. Speech delivered 22 September 2026. <https://www.gov.uk/government/speeches/pm-speech-at-unga-22-september-2026>

⁵⁹ Miliband, E. (23 September 2026). *Foreign Secretary Address to the UNSC on Artificial Intelligence*. Foreign, Commonwealth & Development Office, GOV.UK. <https://www.gov.uk/government/speeches/foreign-secretary-address-to-the-ungsc-on-artificial-intelligence--2>

⁶⁰ Boak, J., and Price, M. L. (14 September 2026). *Trump Dismisses New AI Guardrails, Says There Is a 'SICK Conspiracy' Against AI and Data Centers*. Associated Press. <https://apnews.com/article/trump-ai-guardrails-data-centers-b85df16775ff7e9611a456b061a0e4b9>

⁶¹ United Nations (23 September 2026). *Artificial Intelligence and International Security: Security Council, 10228th Meeting*. UN meeting recording and automatic transcript, statement by Michael Kratsios at 1:13:33. <https://transcripts.un.org/en/sc/10228>

safeguards with China for the development, testing, and use of frontier AI. The senators also called for consideration of an agreement to slow development, potentially including a mutual pause⁶². The letter illustrates differing perspectives within the United States on how international cooperation should complement domestic oversight.

70. Ahead of the 24 September 2026 meeting with Xi Jinping, President of the People's Republic of China, President Trump reiterated support for the existing approach and identified the Department of Justice as the principal safeguard⁶³. This position raises the question of whether established enforcement powers provide sufficient preventive oversight as AI capabilities and applications evolve.
71. China's public position combines support for wider access to AI with opposition to technological exclusion. On 14 September 2026, Foreign Ministry spokesperson Guo Jiakun called for an open and inclusive approach and warned that confrontation would undermine global AI governance. Separately, the state-backed Global Times characterized Dario Amodè's proposals as a means of constraining China's development and maintaining U.S. technological dominance⁶⁴. These distinct statements illustrate why the purpose, scope, and application of safety measures can become contested in a competitive strategic environment.
72. Addressing the Security Council on 23 September 2026, Fu Cong, Permanent Representative of the People's Republic of China to the United Nations, supported a central role for the UN in AI governance, broader participation by developing countries, cross-border cooperation on AI risks, and the retention of ultimate human decision-making authority in military applications. Ambassador Fu also warned that exclusive arrangements could fragment technical standards and deepen the digital divide⁶⁵. These priorities reflect support for multilateral cooperation, with particular emphasis on representation, access, and national sovereignty.

⁶² Office of Senator Ben Ray Lujan (23 September 2026). *Lujan, Van Hollen, Colleagues Push Trump to Negotiate AI Guardrails With Xi*. U.S. Senate. Includes the full text of the senators' letter. <https://www.lujan.senate.gov/newsroom/press-releases/lujan-van-hollen-colleagues-push-trump-to-negotiate-ai-guardrails-with-xi/>

⁶³ Heavey, S. (24 September 2026). *Trump Says He'll Discuss AI With Xi but Wants to 'Leave It Exactly Where It Is'*. Reuters. <https://www.reuters.com/business/media-telecom/trump-says-hell-discuss-ai-with-xi-wants-leave-it-exactly-where-it-is-2026-09-24/>

⁶⁴ Chen, L. (14 September 2026). *China State Newspaper Blasts Anthropic's Calls to Slow AI as 'Cold War' Tactic*. Reuters. <https://www.reuters.com/world/china/china-state-newspaper-blasts-anthropics-calls-slow-ai-cold-war-tactic-2026-09-14/>

⁶⁵ Fu, C. (23 September 2026). *Remarks at the UN Security Council Meeting on Artificial Intelligence and International Security*. Permanent Mission of the People's Republic of China to the United Nations. https://un.china-mission.gov.cn/eng/dbttx_141670/czdbzjds/fcdshd/202609/t20260924_12030908.htm

73. At the White House welcoming ceremony on 24 September 2026, Xi Jinping, President of the People's Republic of China, stated that China and the United States, as leading AI powers, shared the capability and responsibility to develop and manage AI for public benefit and to ensure that AI development remained under human control. President Xi's remarks provide an explicit basis for discussing shared responsibility for AI governance⁶⁶. Effective cooperation would require agreement on applicable safeguards, the evidence to be exchanged, and the institutions responsible for implementation.
74. The United States and China have announced a bilateral dialogue on AI risks and benefits and an agreement to establish a communication channel for AI incidents. The 25 September 2026 White House fact sheet called the forum the U.S.-China "Super Intelligence" Dialogue. China's Ministry of Foreign Affairs confirmed the agreement on 26 September 2026, referring to an AI dialogue, a further exchange in November 2026, and a bilateral channel for AI incidents⁶⁷. A separate Chinese Foreign Ministry statement described "super intelligence" as terminology chosen by the United States⁶⁸. The public statements therefore confirm the announced incident channel but do not yet specify notification thresholds, designated contacts, information-sharing procedures, or response obligations. Those details will determine its practical value during a serious incident.
75. The European Union combines a regulatory framework with policies intended to strengthen investment, innovation, and technological capacity. The AI Act establishes differentiated obligations for providers and deployers, including requirements addressing general-purpose models that present systemic risks⁶⁹. Its influence will depend on effective implementation and access to the expertise needed to assess compliance. For the Union and its partners, regulatory

⁶⁶ Xi, J. (24 September 2026). *Remarks by H.E. Xi Jinping, President of the People's Republic of China, at the Arrival Ceremony at the White House*. Ministry of Foreign Affairs of the People's Republic of China.

https://www.fmprc.gov.cn/eng/gjhdq_665435/3376_665447/3432_664920/3435_664926/202609/t20260924_12030982.html;
RTHK (24 September 2026). *Trump Hails Friendship With Xi at White House Ceremony*. English reporting, AFP/Xinhua.
<https://news.rthk.hk/rthk/en/component/k2/1871461-20260924.htm>

⁶⁷ The White House. (25 September 2026). *Fact Sheet: President Donald J. Trump Advances a Fair and Reciprocal Relationship with China While Hosting Historic State Visit*. <https://www.whitehouse.gov/fact-sheets/2026/09/fact-sheet-president-donald-j-trump-advances-a-fair-and-reciprocal-relationship-with-china-while-hosting-historic-state-visit/>; Ministry of Foreign Affairs of the People's Republic of China. (26 September 2026). *China and the United States Reach Eight Deliverables and Understandings*. https://www.fmprc.gov.cn/eng/xw/zyxw/202609/t20260926_12031663.html

⁶⁸ Ministry of Foreign Affairs of the People's Republic of China. (26 September 2026). *Foreign Ministry Spokesperson's Remarks on Issue Concerning Artificial Intelligence*.
https://www.fmprc.gov.cn/mfa_eng/xw/fyrbt/202609/t20260926_12031649.html

⁶⁹ European Commission (accessed 24 September 2026). *AI Act. Shaping Europe's Digital Future*. <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>

authority must be supported by sufficient technical capacity to evaluate systems and respond to failures.

76. In the State of the Union address on 16 September 2026, Ursula von der Leyen, President of the European Commission, proposed closer cooperation with partners, including Canada and the United Kingdom, on model evaluation, verification, early warning, and AI security. President von der Leyen also announced plans to convene frontier AI laboratories to discuss support for efforts to pace development⁷⁰. Such cooperation could enable participating countries to pool expertise and establish compatible assessment practices. Effective implementation would require sustained access to relevant evidence and clear procedures for responding to evaluation findings.
77. The BRICS New Delhi Declaration: Building for Resilience, Innovation, Cooperation and Sustainability, adopted on 12 September 2026, connects AI governance with development and wider participation. It supports international cooperation to improve access to AI resources, address safety and security, and advance trustworthy applications. It also commits members to implementing the BRICS Leaders' Statement on the Global Governance of Artificial Intelligence⁷¹. This position reinforces the importance of access and representation, particularly for the Global South. It establishes shared political objectives while leaving national authorities to determine their respective regulatory arrangements.
78. Cooperation is also developing through initiatives involving countries across several regions. A Call for Control of Frontier AI Models⁷², published by the government of the Netherlands on 22 September 2026, calls for independent evaluation before deployment, shared reporting of serious incidents, and wider access to scientific expertise. Its signatories include senior representatives of Bahrain, Türkiye, and the United Arab Emirates. The statement also proposes exploring an international institution with functions relating to standards, verification, and consultation when capability thresholds are crossed. This is a political call for

⁷⁰ von der Leyen, U. (16 September 2026). *2026 State of the Union Address by President von der Leyen*. European Commission. Reproduced by the European External Action Service. https://www.eeas.europa.eu/delegations/council-europe/2026-state-union-address-president-von-der-leyen_en

⁷¹ BRICS (12 September 2026). *New Delhi Declaration: Building for Resilience, Innovation, Cooperation and Sustainability*. Prime Minister's Office, Government of India, paragraph 81. https://www.pmindia.gov.in/en/news_updates/brics-new-delhi-declaration-building-for-resilience-innovation-cooperation-and-sustainability-september-12-2026/

⁷² Government of the Netherlands (22 September 2026). *A Call for Control of Frontier AI Models*. Joint statement by international leaders. <https://www.government.nl/documents/2026/09/22/a-call-for-control-of-frontier-ai-models>

further action whose implementation would require subsequent decisions by the participating governments and relevant institutions.

79. Institutional proposals are also emerging below the intergovernmental level. In July 2026, the American Chamber of Commerce in Italy proposed Milan as a possible location for a future international body or platform with a UN mandate on the ethical governance of AI. Lombardy's Open Innovation platform described the initiative as an initial proposal without a formal nomination procedure or agreed timetable⁷³. Its proposed functions include independent risk assessment, greater convergence of standards, and wider participation in international governance. Any institutional role for Milan would require subsequent national support and decisions through the relevant UN processes.
80. Addressing the opening of the UN General Assembly's general debate on 22 September 2026, António Guterres, Secretary-General of the United Nations, placed AI within a broader shift in power toward a small number of private actors and increasingly capable machines. Secretary-General Guterres urged countries leading AI development to exchange information on emerging risks, cooperate on testing, and establish common safeguards. Building on existing UN initiatives, Secretary-General Guterres called for responsible pacing of development and a multilateral AI risk-management framework supported by credible, independent oversight. The address also emphasized the protection of children and the preservation of human authority over decisions involving life and death⁷⁴.
81. At its 10228th meeting on 23 September 2026, the Security Council heard Professor Yoshua Bengio, Co-Chair of the UN Independent International Scientific Panel on AI, and the chief executives of OpenAI, Anthropic, and Hugging Face on AI and international security⁷⁵. Professor Bengio warned that loss of human control could create risks extending across borders and called for independent assessment, incident reporting, and a shared definition of serious

⁷³ Open Innovation, Lombardy Region (28 August 2026). *Milan Proposed as a Global Hub for AI Governance*. <https://openinnovation.regione.lombardia.it/it/news/news/9413/milano-candidata-a-hub-globale-per-la-governance-dell-ia>; Baldini, A. (28 July 2026). *AmCham Italy Proposes Milan as a UN Hub for Artificial Intelligence*. <https://onuitalia.com/2026/07/28/amcham-italy-candida-milano-a-hub-onu-per-lintelligenza-artificiale/>

⁷⁴ Guterres, A. (22 September 2026). *'The Faultlines in Our World Are Widening' – UN Secretary-General at the General Assembly*. United Nations Sustainable Development Group. Full address. <https://unsdg.un.org/latest/announcements/faultlines-our-world-are-widening-%E2%80%93-un-secretary-general-general-assembly>; UN News (22 September 2026). *Guterres: 'Cracks Have Become Canyons' as World Faces Four Tests of Power*. United Nations. <https://news.un.org/en/story/2026/09/1168392>

⁷⁵ United Nations (23 September 2026). *Artificial Intelligence and International Security: Security Council, 10228th Meeting*. UN meeting recording and automatic transcript. <https://transcripts.un.org/en/sc/10228>

incidents⁷⁶. These assessments and proposals were presented by Professor Bengio; the Council concluded the meeting without adopting a new AI governance measure.

82. During the same briefing, OpenAI Chief Executive Officer Sam Altman called for comparable standards on capability assessment, risk evaluation, incident reporting, and human oversight, while leaving the incorporation of such standards into law to national governments. Anthropic Chief Executive Officer Dario Amodei proposed independent access to frontier models, common testing standards for misuse and loss-of-control risks, and notification of incidents significant to global security. Hugging Face Chief Executive Officer Clément Delangue called for greater transparency about agent activity and cyber incidents and argued that defenders need access to capable AI tools⁷⁷. These are positions advanced by the speakers, with different emphases concerning openness and the governance of frontier systems.
83. The Council discussion also showed different views among states. Bahrain called for human oversight in security decisions and referred to its support for an international treaty on AI governance⁷⁸. Russia questioned whether broad AI regulation belonged within the Security Council's mandate and favored inclusive, specialized intergovernmental forums for cyber and military applications⁷⁹. These positions suggest that cooperation on incident reporting or evaluation may be more immediately attainable than agreement on the mandate of a single global institution.
84. On the occasion of the 81st UNGA on 24 September 2026, PAM organized a high-level side event at UNHQ in New York. In his keynote address, H.E. Dr. Mohamed Al Kuwaiti, Head of Cybersecurity for the UAE Government, linked the expansion of digital and agentic systems across essential sectors to the need for security by design. Dr. Al Kuwaiti cited energy and aviation among the sectors undergoing digital transformation and emphasized cooperation among public authorities, the private sector, and communities. The intervention underscores

⁷⁶ Bengio, Y. (23 September 2026). *Professor Yoshua Bengio's Speech at the UN Security Council*. Yoshua Bengio. <https://yoshuabengio.org/en/blog/professor-yoshua-bengios-speech-un-security-council>

⁷⁷ United Nations (23 September 2026). *Artificial Intelligence and International Security: Security Council, 10228th Meeting*. UN meeting recording and automatic transcript, statements by Sam Altman, Dario Amodei, and Clément Delangue. <https://transcripts.un.org/en/sc/10228>

⁷⁸ United Nations (23 September 2026). *Artificial Intelligence and International Security: Security Council, 10228th Meeting*. UN meeting recording and automatic transcript, Bahrain statement at 1:18:24. <https://transcripts.un.org/en/sc/10228>

⁷⁹ Repkin, A. (23 September 2026). *Statement at a UN Security Council Briefing on Artificial Intelligence and International Security*. Permanent Mission of the Russian Federation to the United Nations. <https://russiaun.ru/en/news/223092026>

the importance of incorporating cybersecurity requirements into deployment and strengthening resilience across interconnected services⁸⁰.

85. Human rights provide a further basis for international cooperation. In an open letter dated 14 September 2026, Volker Türk, United Nations High Commissioner for Human Rights, called for stronger regulation, reporting of serious incidents, verification of AI capabilities, and human rights due diligence. High Commissioner Türk emphasized that no country could govern AI alone and that decisions concerning risks with global consequences should not rest with individual companies⁸¹. This approach places accountability to affected individuals and communities at the center of AI governance, including those outside the jurisdictions in which AI systems are developed.
86. At a White House meeting on 29 September 2026, President Donald Trump and executives from Google, Anthropic, Meta, OpenAI, xAI, and NVIDIA signed the White House Accord on Super Intelligence: Joint Commitment on Frontier Responsibilities⁸². The voluntary accord calls for internal controls during model training and deployment, a dedicated team to assess and address failures, evaluation by an independent external auditor or evaluator, and oversight by an independent board committee. The signatories also committed to regular meetings on standards and best practices, while acknowledging that these measures could eventually be codified in law. The published text does not establish enforceable audit standards, public incident-reporting requirements, or a cross-border oversight mechanism. The practical significance of the accord will depend on the independence and scope of external evaluations and on whether identified failures lead to corrective action.
87. Separately, on 29 September 2026, President Donald Trump issued the executive order Inaugurating the Era of Super Intelligence. The order directs federal executive agencies, to the extent permitted by law, to use “Super Intelligence” and “SI” in place of “Artificial Intelligence” and “AI” in future nonstatutory materials. The existing statutory definition of AI remains applicable for implementation of the order, while proposed legislative language is to

⁸⁰ United Nations Web TV. (24 September 2026). *Governance, Technology and Trust in the Age of AI: The Role of Parliaments and Parliamentary Diplomacy in Countering Online Terrorist Narratives*. Intervention by H.E. Dr. Mohamed Al Kuwaiti, 26:25–28:28. <https://webtv.un.org/en/asset/k1a/k1a7o2u5ue>

⁸¹ Steinberg, I. (14 September 2026). *Countries Must Increase AI Regulation to Avoid ‘Existential Risks’*: Türk. UN News, reporting on the High Commissioner’s open letter. <https://news.un.org/en/story/2026/09/1168326>

⁸² Rapid Response 47. (29 September 2026). White House Accord on Super Intelligence: Joint Commitment on Frontier Responsibilities. Official White House Rapid Response account, X. <https://x.com/RapidResponse47/status/2105052408032403532>

be submitted within 60 days⁸³. The measure changes federal terminology and initiates consideration of a legal definition; it does not itself establish additional safety obligations for developers.

88. Taken together, the positions examined in this section demonstrate broad recognition that advanced AI requires safeguards, but continuing disagreement over the authority to assess risks and require corrective action. The United Kingdom and European Union have emphasized evaluation and international cooperation, with the European Union also pursuing binding regulation. The United States has combined reliance on domestic authority and technological leadership with voluntary commitments by leading developers. China has called for human control, wider participation, and opposition to technological exclusion. These differences become consequential when a serious incident requires independent access to evidence, timely intervention, and cooperation across borders. This comparison reflects the positions examined here and does not imply uniform views within any government, industry, or national research community.
89. These developments raise questions about institutional authority as well as the content of safeguards. Companies can provide essential technical evidence and undertake immediate interventions, but decisions concerning acceptable public risk require lawful authority and independent scrutiny. Governments, in turn, need arrangements that permit qualified evaluators to examine relevant systems and communicate findings to bodies empowered to act. Confidentiality protections should secure sensitive information while allowing the access necessary for effective oversight. Separately, on 29 September 2026, President Trump signed an Executive Order establishing America.gov as a unified digital portal for federal services. The order requires “super intelligence” used in connection with the platform to be accurate, reliable, and transparent, and establishes safeguards including data minimization, secure authentication, and auditable authorization. The initiative provides an example of the Administration applying operational safeguards to the deployment of advanced AI within federal public services⁸⁴.

⁸³ The White House. (29 September 2026). Inaugurating the Era of Super Intelligence. <https://www.whitehouse.gov/presidential-actions/2026/09/inaugurating-the-era-of-super-intelligence/>

⁸⁴ The White House. (29 September 2026). Streamlining Access to Government Services Through America.gov. Presidential Actions. <https://www.whitehouse.gov/presidential-actions/2026/09/streamlining-access-to-government-services-through-america-gov/>

90. Strategic competition may impede international cooperation on AI governance. Governments may view restraint as creating a disadvantage relative to competitors that continue advancing AI capabilities. Requirements for evaluation or disclosure may also face resistance when oversight institutions are associated with rival jurisdictions. Measures addressing identifiable risks should therefore rest on clear technical evidence, remain proportionate to those risks, and apply under transparent conditions. Credibility depends on consistent implementation and the ability to assess compliance.
91. The distribution of AI activities across jurisdictions can also complicate enforcement. Development, computing infrastructure, integration, and deployment may be controlled by different organizations in different countries. An investigation may consequently depend on several legal authorities and access to records held abroad. Compatible requirements for preserving evidence and reporting serious incidents would reduce avoidable delays. Regulatory competition that weakens these requirements could transfer risks to countries with little influence over the activities that generate them.
92. Technological dependence adds a further constraint. The OECD's Competition in Artificial Intelligence Infrastructure examines competition concerns arising from the computing resources and supply chains on which AI development depends⁸⁵. For public institutions, reliance on external infrastructure and service providers can affect access to operational information, security updates, and continuity arrangements. Procurement and partnership decisions therefore have implications for a state's future capacity to assess risks and maintain essential services.
93. Wider access to models can support research, local adaptation, and reduced dependence on individual providers. The OECD's AI Openness: A Primer for Policymakers distinguishes different forms of openness and examines the benefits and risks of releasing model weights⁸⁶. Locally operated systems may give deploying institutions greater control over their data while limiting the original provider's visibility and ability to intervene remotely. These arrangements require responsibilities suited to the actors that actually operate, modify, and secure the systems concerned.

⁸⁵ Organisation for Economic Co-operation and Development (OECD) (14 November 2025). *Competition in Artificial Intelligence Infrastructure*. OECD Roundtables on Competition Policy Papers, No. 330. OECD Publishing. https://www.oecd.org/en/publications/competition-in-artificial-intelligence-infrastructure_623d1874-en.html

⁸⁶ Organisation for Economic Co-operation and Development (OECD) (14 August 2025). *AI Openness: A Primer for Policymakers*. OECD Artificial Intelligence Papers, No. 44. OECD Publishing. <https://doi.org/10.1787/02f73362-en>

94. For the Euro-Mediterranean and Gulf regions, the practical implications concern the security of essential services, access to expertise, and the terms of technological cooperation. National assessments should consider dependencies in energy, transport, telecommunications, finance, and public administration, including disruption originating outside the jurisdiction. Regional cooperation could help institutions share expertise, compare procurement requirements, and establish reliable channels for incident response while respecting different constitutional and regulatory systems.
95. The central geopolitical question is whether states can establish credible mutual assurance while continuing to compete. Shared commitments become operationally meaningful when agreements specify the evidence required, the authorities responsible for reviewing that evidence, and the measures to be taken when safeguards fail. Progress will depend on participation beyond the leading AI-producing states and on arrangements that uphold human rights, support national implementation, and enable cooperation on shared risks despite broader political differences.

The Current State of AI Governance and Prospects for Cooperation

96. By late September 2026, international cooperation on AI-related risks had gained a dedicated United Nations dialogue and a declaration addressing digital crime, while operational arrangements for independent oversight of frontier AI remained at an early stage. The first United Nations Global Dialogue on AI Governance convened in Geneva in July 2026, with a second session scheduled for New York in May 2027⁸⁷. On 26 September 2026, the 15th United Nations Congress on Crime Prevention and Criminal Justice adopted the Abu Dhabi Declaration, addressing the criminal exploitation of digital technologies and safeguards for the use of technology in criminal justice, including human oversight, transparency, due process, and data protection⁸⁸. At the close of the General Debate on 28 September 2026, the President of the General Assembly reported that 128 delegations had addressed AI and called for human-centered, inclusive governance attentive to the needs of developing countries⁸⁹. The Dialogue provides an inclusive forum for cooperation, while the Declaration addresses crime prevention

⁸⁷ United Nations. (n.d.). Global Dialogue on AI Governance. <https://www.un.org/global-dialogue-ai-governance/en>

⁸⁸ UN News. (26 September 2026). Countries adopt global blueprint to tackle fast-evolving crime in digital age. <https://news.un.org/en/story/2026/09/1168455>

⁸⁹ United Nations General Assembly. (28 September 2026). Remarks at the close of the general debate of the eighty-first session. <https://www.un.org/pga/81/documents/speeches/close-of-the-general-debate-28-september-2026/>

and criminal justice. Neither development, by itself, establishes a system for independently verifying frontier AI safeguards or coordinating responses to serious incidents across borders.

97. Recent public interventions have sharpened the distinction between two risk pathways examined in this report. At Colgate University on 18 September 2026, former U.S. President Barack Obama warned about both potential misalignment in increasingly capable systems and the malicious use of AI by human actors. President Obama characterized human-directed misuse as the more immediate concern and argued that voluntary restraint could provide only an interim response pending effective government regulation⁹⁰. In an interview on 27 September 2026, Bill Gates similarly emphasized harmful use by human actors and called for legally required monitoring and records of consequential AI activity. Mr. Gates argued that an emergency shutdown mechanism alone would be insufficient without the information needed to identify and investigate harmful conduct⁹¹. These interventions are assessments and policy proposals; the relative prevalence of either risk pathway must be determined from incident evidence.
98. The announced U.S.-China dialogue creates a potential channel for discussing AI risks despite wider strategic competition. The practical test will be whether both governments specify the types of incidents requiring communication, the authorities responsible for notification, and the information that can be exchanged promptly and securely. The public statements examined above do not yet establish those operational details⁹². Bilateral arrangements could contribute to risk reduction, but incidents affecting other countries will also require channels through which relevant national authorities can obtain timely information and technical assistance.
99. The development of a national AI model does not, by itself, confer control over every part of the technological system on which that model depends. An application adapted from an

⁹⁰ Colgate University. (18 September 2026). *A Conversation with Barack Obama*, full event recording. <https://www.youtube.com/watch?v=mgfO4nnPL3Y>. Relevant remarks are also available in the [recorded excerpt provided for this report](#), approximately 5:40–6:45 and 12:26–14:56.

⁹¹ NBC News. (27 September 2026). *Bill Gates Says It's "Not Enough to Have a Kill Switch" for AI: Full Interview*. <https://www.nbcnews.com/meet-the-press/video/bill-gates-says-it-s-not-enough-to-have-a-kill-switch-for-ai-full-interview-270615621979>. See also Marquez, A. (27 September 2026). *Bill Gates Says Getting Countries to Agree on AI Regulations Will Be Harder Than Cold War-Era Nuclear Negotiations*. NBC News, published by NBC4 Washington. <https://www.nbcwashington.com/news/national-international/bill-gates-getting-countries-to-agree-on-ai-regulations-harder-than-cold-war-era-nuclear-negotiations/4159873/>

⁹² The White House. (25 September 2026). *Fact Sheet: President Donald J. Trump Advances a Fair and Reciprocal Relationship with China While Hosting Historic State Visit*. <https://www.whitehouse.gov/fact-sheets/2026/09/fact-sheet-president-donald-j-trump-advances-a-fair-and-reciprocal-relationship-with-china-while-hosting-historic-state-visit/>; Ministry of Foreign Affairs of the People's Republic of China. (25 September 2026). *President Xi Jinping Holds Talks with U.S. President Donald J. Trump*. https://www.fmprc.gov.cn/mfa_eng/xw/zyxw/202609/t20260925_12031181.html

existing foundation model may carry forward limitations of the underlying model; Stanford University's research on foundation models identifies the transmission of defects to downstream adaptations as a consequence of their widespread reuse⁹³. A separately trained model may likewise draw on earlier published architectures and depend on external computing infrastructure. The UAE's Falcon-40B illustrates this distinction: its model card describes an architecture adapted from published GPT-3 research, a custom training codebase, and training on AWS cloud infrastructure⁹⁴. The OECD documents the wider supply chain linking AI development to chips, data centers, and cloud providers⁹⁵. As AI becomes embedded in essential services, national capacity to develop a model must therefore be accompanied by the capacity to assess its dependencies, investigate failures, and maintain effective control over its deployment.

100. This distribution of technological dependence and exposure makes oversight a concern for countries beyond the principal AI-producing states. At the South by Southwest (SXSW) conference on 11 March 2018, Elon Musk argued that AI posed a greater danger than nuclear warheads and called for a public body with insight into, and oversight of, its development⁹⁶. OpenAI subsequently proposed consideration of an authority analogous to the International Atomic Energy Agency for superintelligence projects exceeding defined capability or computing thresholds⁹⁷. The comparison concerns independent verification and accountable control of potentially severe risks; the technologies do not have identical mechanisms of harm. A negotiated framework, implemented through national law in countries hosting leading developers, could provide qualified independent evaluators with access to high-risk systems, establish thresholds for serious-incident notification, and support coordinated corrective action. Countries that depend on these systems should also participate meaningfully in setting common safeguards and obtaining information when incidents threaten public services. The

⁹³ Bommasani, R., et al. (2021). *On the Opportunities and Risks of Foundation Models*. Stanford Center for Research on Foundation Models. <https://crfm.stanford.edu/report>

⁹⁴ Technology Innovation Institute. (2023). *Model Card for Falcon-40B*. <https://huggingface.co/tiiuae/falcon-40b/blob/main/README.md>

⁹⁵ Organisation for Economic Co-operation and Development (OECD). (2025). *Competition in Artificial Intelligence Infrastructure*. https://www.oecd.org/en/publications/competition-in-artificial-intelligence-infrastructure_623d1874-en/full-report.html

⁹⁶ SXSW. (11 March 2018). *Elon Musk Answers Your Questions! | SXSW 2018*. Event recording. <https://www.youtube.com/watch?v=kzIUyrcb0s>

⁹⁷ OpenAI. (22 May 2023). *Governance of Superintelligence*. <https://openai.com/index/governance-of-superintelligence/>

United Nations Global Dialogue on AI Governance provides an existing inclusive forum for that discussion⁹⁸.

101. The Euro-Mediterranean and Gulf regions have an institutional basis for more systematic observation of AI-related risks. PAM-CGS operates a parliamentary observatory on the misuse of AI and information and communication technologies and publishes a digest on AI and emerging technologies⁹⁹. On 23 and 24 June 2026, the Council of Europe brought together authorities responsible for AI governance and data protection from the Southern Mediterranean and Europe to examine practical challenges and exchange experience¹⁰⁰. At the PAM side event on 24 September 2026 at UNHQ in New York, Hon. Giulio Centemero, President of PAM, proposed establishing an AI Observatory for the Wider Mediterranean¹⁰¹. The proposed regional forum has not yet been established and is distinct from the existing PAM-CGS observatory. Existing mechanisms provide a basis for monitoring documented incidents, legislative developments, institutional capacity, and risks to essential services. If established, the proposed forum could extend this work through wider regional cooperation.

102. The next phase should convert general commitments into procedures that can be tested. Priority questions concern access to evaluation evidence, common thresholds for serious-incident reporting, preservation of records, cross-border notification, and the authority to require corrective action. Regional monitoring can help identify where these procedures are absent or incompatible, while parliamentary review can determine whether national institutions possess the mandate and resources necessary to implement them.

Recommendations for Parliamentary Action in the Euro-Mediterranean and Gulf Regions

103. Parliamentary action should connect legal obligations with the institutional capacity required for implementation. Priorities for the Euro-Mediterranean and Gulf regions include clear responsibility for high-risk AI deployments, access to reliable evidence, and effective powers to intervene when serious risks arise. The Parliamentary Assembly of the

⁹⁸ United Nations. (2026). *Global Dialogue on AI Governance*. <https://www.un.org/global-dialogue-ai-governance/en>

⁹⁹ Parliamentary Assembly of the Mediterranean (PAM). (18 June 2026). PAM contributes to OECD discussions on AI governance and international cooperation. <https://pam.int/pam-contributes-to-oecd-discussions-on-ai-governance-and-international-cooperation/>; PAM-CGS. (n.d.). PAM-CGS Digest on AI and Emerging Technologies. <https://www.cgspam.org/digest/>

¹⁰⁰ Council of Europe. (n.d.). Advancing data protection and AI governance in the Southern Mediterranean region. <https://pjp-eu.coe.int/en/web/south-programme/-/advancing-data-protection-and-ai-governance-in-the-southern-mediterranean-region>

¹⁰¹ Dal Maso, E. (29 September 2026). AI, Centemero: An Observatory in the Mediterranean and Milan as a Candidate for a Future UN Body [translated title]. MilanoFinanza. <https://www.milanofinanza.it/news/ai-centemero-un-osservatorio-nel-mediterraneo-e-milano-candidata-a-una-futura-struttura-onu-202609291151081667>

Mediterranean (PAM) can support national parliaments through comparative legislative work, dialogue with competent authorities, and cooperation on issues with cross-border implications. The following recommendations should be adapted to national constitutional arrangements and the risks associated with specific AI applications.

104. At the PAM High-Level Side Event, *Governance, Technology and Trust in the Age of AI: The Role of Parliaments and Parliamentary Diplomacy in Countering Online Terrorist Narratives*, held on 24 September 2026 at UNHQ in New York, H.E. Hon. Mohamed Abou El Enein (Egypt), PAM President Emeritus, emphasized the need for immediate action in response to the growing gap between the speed of AI development and the pace of the international response. Hon. Abou El Enein recognized the potential benefits of AI in information, communication, and production, while expressing concern about risks to humanity in the context of the substantial profits generated by AI companies. Hon. Abou El Enein called for parliaments and governments to work together on legislation, with PAM playing an important role, and urged the prompt dissemination of shared rules to address these challenges¹⁰².
105. Coordinated legislative work could build on existing initiatives. The United Nations Office of Counter-Terrorism (UNOCT) report on the February 2026 parliamentary training in Manila describes plans to prepare baseline legislative recommendations and draft model provisions on AI in terrorism and counterterrorism, in consultation with the Office of the United Nations High Commissioner for Human Rights (OHCHR) and participating parliamentary institutions¹⁰³. These resources could inform national reviews within the scope of counterterrorism. Broader AI legislation would also require consultation with relevant sectoral authorities, scientific institutions, civil society, and affected communities.

¹⁰² United Nations Web TV. (24 September 2026). *Governance, Technology and Trust in the Age of AI: The Role of Parliaments and Parliamentary Diplomacy in Countering Online Terrorist Narratives*. <https://webtv.un.org/en/asset/k1a/k1a7o2u5ue>

¹⁰³ United Nations Office of Counter-Terrorism (UNOCT). (16 April 2026). *Parliamentary AIPA Manila Report, February 2026*. https://www.un.org/counterterrorism/sites/default/files/2026-04/parliamentary_aipa_manila_report_february_2026.pdf

Figure 2: From AI Risk to Institutional Response: A Parliamentary Framework for Governance, Oversight and Regional Cooperation¹⁰⁴



Legislative Priorities

106. Parliaments should assess whether existing legislation adequately addresses the capabilities and operating conditions of high-risk AI systems. Legislative reviews should identify specific gaps in responsibility, investigative powers, and remedies before additional obligations are introduced. Regulatory requirements should reflect the severity and likelihood of harm, the scale of deployment, the authority delegated to the system, and the reversibility of consequential actions.
107. Legal definitions should distinguish the underlying AI model from the agent or application operating through that model. Relevant functions include the use of external tools, access to credentials, communication with other systems, and the pursuit of objectives through multiple steps. Common terminology for serious incidents, unauthorized behavior, and effective human

¹⁰⁴ Graphic compiled by the CGS research team, based on the analysis presented in this report.

control would support cooperation across jurisdictions. PAM could facilitate comparison of national definitions and the preparation of explanatory guidance for parliamentary committees.

108. Responsibility should correspond to the role and effective control of each actor. Legislation should clarify the duties of developers, deploying organizations, infrastructure operators, and users, including when an incident involves failures at several levels. Contractual arrangements should allocate responsibility for security updates, subcontractors, access to evidence, and cooperation with investigations. Public institutions should remain accountable for deployment decisions, and outsourcing should not diminish statutory duties.
109. High-risk AI deployments should undergo documented assessment before use, after material changes, and at intervals proportionate to the risk. Assessments should examine the complete application and operating environment, including permissions, external inputs, and connected services. The National Institute of Standards and Technology (NIST) Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile provides a reference for considering risks throughout the system life cycle¹⁰⁵. When potential consequences are serious, evaluations should include independent expertise and clearly defined conditions for corrective action, restricted deployment, or suspension.
110. Legislation should require timely reporting of serious AI incidents to a designated competent authority. Reporting thresholds should address harm and material risk, including significant loss of control, unauthorized access to sensitive systems, and failures of containment. Initial notifications should provide available information without awaiting a completed investigation, with subsequent updates reflecting emerging findings. Procedures should also provide for notification of affected institutions and foreign authorities when consequences may extend across borders.
111. Incident reports should distinguish immediate interruption from effective containment. Reports should identify affected resources, response measures, and any remaining access or capability that could allow harmful activity to continue. Evidence-preservation requirements should cover the records necessary to reconstruct consequential actions while respecting data protection, due process, and legitimate confidentiality. Investigators should be able to examine

¹⁰⁵ National Institute of Standards and Technology (NIST). (July 2024). *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile* (NIST AI 600-1). U.S. Department of Commerce. <https://doi.org/10.6028/NIST.AI.600-1>

human instructions, external manipulation, and agent behavior together when determining the origin and progression of an incident.

112. Meaningful human control should be required for consequential decisions. Designated personnel must have the information, time, authority, and practical means necessary to assess and interrupt an action. Sectoral authorities should specify appropriate levels of supervision for critical infrastructure, health care, law enforcement, border management, and other sensitive applications. Responsibilities should remain clearly allocated when several agents or organizations participate in the same process.
113. Defense applications, weapons-related activity, and dual-use biological research require specialized legal, scientific, and security review. Parliaments should examine whether existing oversight can identify connected activities that present serious risks even when individual requests appear legitimate. Measures should address relevant pathways to harm while preserving lawful research, public health preparedness, and scientific cooperation. The International Committee of the Red Cross (ICRC) publication, *The Rise of Artificial Intelligence Requires a Human-Centred Approach in Conflict Zones*¹⁰⁶, provides a relevant foundation for preserving human judgment and responsibility in military applications.
114. Legal obligations should remain applicable when AI systems are modified, hosted locally, operated offline, or distributed across multiple infrastructures. In these circumstances, legislation should identify the organizations responsible for access controls, assessment, evidence preservation, and incident response. Requirements should reflect capability and deployment risk. Open access should neither provide an exemption from accountability nor justify restrictions unrelated to a demonstrated risk.
115. Legislation addressing AI-enabled manipulation and terrorist recruitment should be accompanied by prevention measures protecting children and other vulnerable users. Parliaments should examine age-appropriate design, accessible reporting channels, digital literacy, and support for educators and families. Parliamentary scrutiny should also address the use of AI by public authorities for these purposes, including accuracy, discriminatory effects,

¹⁰⁶ International Committee of the Red Cross (ICRC). (24 July 2023). *The Rise of Artificial Intelligence Requires a Human-Centred Approach in Conflict Zones*. <https://www.icrc.org/en/document/rise-artificial-intelligence-requires-human-centred-approach-conflict-zones>

and access to review. Security objectives should be pursued through precise, lawful measures protecting legitimate expression and civic participation.

Technical and Institutional Priorities

116. Institutional readiness should be assessed before a public authority commits to AI deployment. At the PAM side event held at UNHQ on 24 September 2026, Alexandre Zouev, Acting Under-Secretary-General for Counter-Terrorism at the United Nations Office of Counter-Terrorism (UNOCT), emphasized that responsible assessment may require redesigning, postponing, or declining an AI application, and linked effective oversight to adequate authority, expertise, and resources¹⁰⁷. Procurement approval and budget scrutiny should therefore establish whether the deploying institution can supervise the proposed system throughout operation and respond to foreseeable failures.
117. Public procurement requirements should specify the evidence required from suppliers concerning AI capabilities, limitations, evaluation, and security. Contracts should provide appropriate access to operational records, cooperation during serious incidents, and clear responsibility for updates and subcontractors. Assessments for essential services should also address provider withdrawal, service interruption, changes in contractual terms, and the capacity to transfer services to another provider or restore human-operated processes.
118. Technical controls should limit each agent to the resources and authority required for the assigned task. Appropriate measures include restricted network access, secure credential management, isolated operating environments, and additional human authorization for consequential actions. In *Managing the Cyber Risk of Agentic AI*¹⁰⁸, the United Kingdom's National Cyber Security Centre (NCSC) recommends controls across the broader agent operating environment, including accessible systems, data, and credentials. Deploying institutions should test these controls against the actual configuration in use.
119. Institutions should maintain protected records sufficient to reconstruct consequential activity. Records should identify assigned objectives, material instructions and external inputs, permissions, tool use, communications, and interventions. Monitoring should trigger review of

¹⁰⁷ United Nations Web TV. (24 September 2026). *Governance, Technology and Trust in the Age of AI: The Role of Parliaments and Parliamentary Diplomacy in Countering Online Terrorist Narratives*. UNOCT intervention, 1:34:32–1:35:44. <https://webtv.un.org/en/asset/k1a/k1a7o2u5ue>

¹⁰⁸ National Cyber Security Centre (NCSC). (20 August 2026). *Managing the Cyber Risk of Agentic AI*. Sections 3 and 4, on oversight, sandboxing, network access, and credentials. <https://www.ncsc.gov.uk/blogs/managing-the-cyber-risk-of-agentic-ai>

attempts to exceed permissions, conceal activity, or evade termination. Retention periods should be proportionate, access should be restricted, and investigative records should be protected against alteration or deletion.

120. Emergency procedures should enable the suspension of execution, revocation of credentials, isolation of affected resources, and interruption of communications without requiring cooperation from the system under investigation. Tests should cover the complete operating environment and interactions between agents. NCSC guidance emphasizes that halting an agent may require restrictions on infrastructure and communications in addition to termination of the immediate process¹⁰⁹. Response plans should designate the authorities empowered to approve each intervention and specify arrangements for maintaining essential services.
121. Oversight bodies require expertise in AI, cybersecurity, law, human rights, and relevant sectors of deployment. Parliamentary budget scrutiny should assess staffing, resources, and access to independent advice. Countries with limited capacity could benefit from joint training, expert networks, and cooperation with qualified regional or international institutions. These arrangements should strengthen national decision-making and provide safeguards for sensitive information and conflicts of interest.
122. Controlled testing environments can support the assessment of AI applications and refinement of requirements before wider deployment. Testing should operate under defined objectives, supervision, and conditions for termination, preserve applicable protections, and produce findings relevant to regulatory decisions. Periodic exercises should assess coordination among developers, deploying institutions, infrastructure operators, and public authorities during incidents affecting multiple sectors or jurisdictions.
123. Existing cybersecurity and emergency-response structures should perform these functions when the necessary mandate and capacity are in place. Governments should designate contact points responsible for receiving urgent AI incident notifications, obtaining technical assistance, and communicating with foreign counterparts. Parliamentary scrutiny should assess the adequacy of authority and resources for incidents combining cyber disruption, unauthorized agent behavior, and potential physical harm.

¹⁰⁹ National Cyber Security Centre (NCSC). (20 August 2026). *Managing the Cyber Risk of Agentic AI*. Section 7, Emergency Shutdown. <https://www.ncsc.gov.uk/blogs/managing-the-cyber-risk-of-agentic-ai>

Parliamentary Oversight and the Role of PAM

124. Parliaments should require regular reporting on high-risk AI deployments in public institutions and essential services. Committee scrutiny should assess the practical effectiveness of safeguards through evidence on detection and response times, recurring failures, unresolved vulnerabilities, and implementation of corrective measures. When necessary, committees should have access to independent technical advice and procedures for examining sensitive material under appropriate confidentiality protections.
125. Emergency powers should have a clear legal basis, defined thresholds, and proportionate scope. Legislation should designate the authorities empowered to suspend a service, restrict access, or require disclosure of technical information, with applicable procedural safeguards and review mechanisms. Temporary powers should be time-limited and subject to renewed legislative scrutiny before extension. Monitoring and information-sharing should address identifiable risks while protecting privacy, freedom of expression, and due process.
126. Parliamentary institutions should establish rules governing the use of AI in parliamentary work. Discussions at the 24 September 2026 PAM side event at UNHQ addressed work on a charter for parliamentarians and parliamentary staff, including concerns about fabricated references in AI-assisted drafting¹¹⁰, the so-called AI Hallucinations. Internal rules should require verification of legal authorities, quotations, statistics, and sources before official use, protection of confidential material, and clearly assigned human responsibility for final texts. Staff training and procedures for correcting errors should support implementation.
127. The Inter-Parliamentary Union (IPU) Guidelines for AI in Parliaments¹¹¹ provide an existing reference for institutional strategy, governance, ethical assessment, and risk management. Each parliament could adapt the guidance to national parliamentary procedures, define approved uses and tools, and establish requirements for disclosing AI assistance. Parliamentary deliberation and the exercise of statutory powers should remain the responsibility of members and competent parliamentary authorities.

¹¹⁰ United Nations Web TV. (24 September 2026). *Governance, Technology and Trust in the Age of AI: The Role of Parliaments and Parliamentary Diplomacy in Countering Online Terrorist Narratives*. Intervention by Sergio Piazzi, beginning at 1:38:25. <https://webtv.un.org/en/asset/k1a/k1a7o2u5ue>

¹¹¹ Inter-Parliamentary Union (IPU). (2024). *Guidelines for AI in Parliaments*. Published in partnership with the Parliamentary Data Science Hub in the IPU Centre for Innovation in Parliament. <https://www.ipu.org/resources/publications/reference/2024-12/guidelines-ai-in-parliaments>

128. PAM could use its existing PAM-CGS Parliamentary Observatory to develop a dedicated Euro-Mediterranean and Gulf monitoring strand¹¹². Periodic briefs could distinguish verified incidents from provider statements and unresolved reports, compare relevant national legislation, and identify risks affecting interconnected public services. Findings could inform parliamentary hearings, legislative reviews, and regional preparedness exercises. Any exchange of nonpublic incident information should operate through competent national authorities and applicable safeguards for confidentiality, personal data, and due process.
129. PAM could facilitate regular exchanges among parliamentary committees responsible for technology, justice, defense, economic affairs, and human rights. Initial cooperation could address model legislative provisions, terminology for serious incidents, notification procedures, and evidence-preservation requirements. Shared procurement clauses and evaluation criteria could assist parliamentary administrations with limited negotiating or technical capacity. Cooperation should improve practical compatibility among national approaches and support informed legislative decisions.
130. Parliamentary diplomacy should support the participation of countries with limited resources in international AI governance discussions. PAM could bring regional experience into dialogue with the United Nations, the IPU, other international parliaments, relevant regulators, scientific institutions, and industry. Participation should include independent researchers, civil society, and affected communities, with transparent arrangements for managing conflicts of interest. External expertise should support the independent judgment of parliamentary institutions.
131. Implementation should begin with a national review of responsibilities and institutional capacity, followed by priority measures on high-risk procurement, incident reporting, and parliamentary use of AI. Governments should report periodically on progress and risks beyond the scope of existing controls. Detailed technical requirements may require frequent revision, while changes affecting rights, duties, and public powers should remain subject to parliamentary deliberation. PAM can support implementation through sustained comparative legislative exchanges and cooperation on regional preparedness.

¹¹² Parliamentary Assembly of the Mediterranean (PAM). (18 June 2026). *PAM contributes to OECD discussions on AI governance and international cooperation*. <https://pam.int/pam-contributes-to-oecd-discussions-on-ai-governance-and-international-cooperation/>

Conclusions

132. The cases examined in this report indicate that AI-related risks extend beyond the misuse of individual capabilities. Models and agents can interact with external tools, digital services, and infrastructure through connected workflows. AI continues to facilitate established threats, including fraud, impersonation, cyber intrusion, and malicious technical research. Agentic systems introduce an additional concern: harmful outcomes may arise from unauthorized methods adopted in pursuit of an assigned task.
133. The distinction between human-directed misuse and unauthorized agent behavior remains essential, including when the two pathways overlap. An external actor may manipulate a system, delegate part of an operation to an agent, or influence agent behavior through content encountered during a task. Subsequent actions may exceed the scope of the original request or authorization. Investigations should therefore examine the origin of the objective, assigned instructions and permissions, external inputs, tool use, and opportunities for intervention.
134. These developments place the relationship between capability and control at the center of the governance debate. Model-level safeguards remain necessary, but may be insufficient when a system can use credentials, communicate externally, or act through connected services. Access restrictions, monitoring of consequential activity, reliable records, and mechanisms for interrupting execution and communication should form an integrated protective framework.
135. AI providers retain an important role through activity monitoring, incident investigation, and the withdrawal of access. Provider capacity to intervene may be limited when systems operate locally, undergo modification by users, or rely on infrastructure controlled by other organizations. Responsibilities should therefore correspond to the authority exercised by developers, deploying institutions, and infrastructure operators at each stage of an operation.
136. The practical outcome of an intervention requires careful assessment. Suspending an account or blocking a request may interrupt harmful activity without removing previously acquired knowledge, software, or technical material. Assessment should consider the speed of detection, the continuation of related activity through other channels, and the capabilities remaining after access has been withdrawn.
137. This consideration is particularly significant when digital assistance may contribute to physical harm. Cyber operations, critical infrastructure, weapons-related activity, and dual-use biological research require assessment of connected patterns of activity alongside individual

requests. Safeguards should address serious risks while avoiding unnecessary restrictions on legitimate research, access to information, and beneficial uses of AI.


138. The international environment complicates the development of common safeguards. The national and regional positions examined in this report assign differing priorities to safety, competitiveness, sovereignty, and access to technology. Such differences may encourage regulatory competition or lead to perceptions of safety measures as instruments of technological advantage. Strategic tensions may also reduce willingness to exchange information or accept evaluations conducted by institutions associated with competing jurisdictions.
139. Agreement on a single global regulatory model may remain difficult in the immediate future. Cooperation could nevertheless advance through practical requirements concerning definitions of serious incidents, evaluation of high-risk systems, preservation of evidence, and cross-border notification. Agreement on these matters would allow distinct national legislative approaches while strengthening coordinated responses to incidents involving systems, providers, and infrastructure across multiple jurisdictions.
140. Euro-Mediterranean and Gulf countries have a direct interest in such cooperation. Public institutions and economies across these regions are interconnected through energy, maritime transport, telecommunications, finance, cloud services, and digital infrastructure, while technical and regulatory capacities vary considerably. An incident affecting one part of this environment may have consequences elsewhere, including in countries with limited influence over the development or deployment of the system concerned.
141. Effective national oversight does not require every state to reproduce the technical capabilities of frontier developers. National authorities nevertheless need sufficient capacity to assess high-risk deployments, protect essential infrastructure, receive and evaluate warnings, preserve evidence, and coordinate intervention. Regional cooperation can reduce duplication and provide smaller administrations with access to expertise that may be difficult to maintain independently.
142. Parliaments have a distinct responsibility for ensuring that AI governance rests on a clear legal basis and accountable public authority. Legislative action determines the duties of relevant actors, the powers of oversight bodies, and the safeguards governing monitoring, information sharing, and emergency intervention. Parliamentary scrutiny should assess the

practical effectiveness of these powers and ensure proportionate application subject to review. Parliamentary institutions should also govern the use of AI in legislative work through source verification, protection of confidential information, staff training, and clearly assigned human responsibility for official texts.


143. PAM can support this work through sustained dialogue among parliamentarians representing different legal systems, technological capacities, and international partnerships. PAM's contribution can include identifying areas requiring compatibility among national approaches, bringing regional concerns into international discussions, and supporting the participation of Member States with limited regulatory or technical resources.
144. Effective human control requires practical institutional capacity. The presence of a human decision-maker offers limited assurance without adequate information, sufficient time, and the authority and technical means to intervene. Meaningful control depends on clearly assigned responsibility, traceable activity, and the capacity to restrict, suspend, or review consequential actions.
145. AI can support scientific research, economic development, public services, and institutional preparedness across the Euro-Mediterranean and Gulf regions. The central question concerns the capacity of governance arrangements to keep pace with technological development. Sustainable innovation requires institutions capable of assessing risks, intervening when necessary, and ensuring that security measures comply with applicable law and human rights obligations.
146. The evidence examined in this report supports the coordinated development of technical safeguards, national legislation, independent evaluation, and regional and international cooperation. Increasing operational capacity should be accompanied by clearer responsibilities and stronger accountability. For PAM and Member States, timely action can help prevent individual failures from escalating into wider threats to regional security, democratic accountability, and institutional resilience.

Prepared by
the Parliamentary Assembly of the Mediterranean (PAM),
Centre for Global Studies (CGS)

Viale Antonio Onofri, 47890 - 103
Città di San Marino, San Marino



PAM



CGS

www.pam.int
www.cgspam.org

