



# Amazon: G7 Hiroshima AI Process (HAIP) Transparency Report

<https://www.aboutamazon.com/>

Publication date: Nov 4, 2025, 05:27 AM PST

Reporting period: 2025



## Section 1 - Risk identification and evaluation

a. How does your organization define and/or classify different types of risks related to AI, such as unreasonable risks?

Amazon is both a developer and deployer of AI. AWS offers AI services made available through our cloud offerings. While Amazon builds and uses AI technology across many of its businesses, this report focuses on the Amazon Nova models. The Amazon Nova family of models deliver frontier intelligence and industry leading price performance, available in Amazon Bedrock. Where relevant, we include examples of how Amazon's businesses use mechanisms to responsibly develop and deploy AI technology to solve customer needs and advance sustainability goals.

Amazon defines and classifies AI risks in Amazon Nova models through our Frontier Model Safety Framework (FMSF), which establishes Critical Capability Thresholds across key risk categories. We evaluate these models against these thresholds using maximal capability assessments and implement appropriate safeguards before deploying any model that reaches these thresholds.

Please see the "Critical Capability Thresholds" section in our [Frontier Model Safety Framework](#).

**b. What practices does your organization use to identify and evaluate risks such as vulnerabilities, incidents, emerging risks and misuse, throughout the AI lifecycle?**

Our approach includes technical safeguards, policy controls, and continuous evaluation processes designed to identify and address frontier model vulnerabilities while promoting safe and responsible AI development. Amazon’s evaluation methodology combines automated benchmarking, expert red-teaming, and human-centric risk assessments to identify and evaluate risks throughout a frontier model’s lifecycle. These model evaluations are conducted on an ongoing basis, including during training and prior to deployment of new frontier models. Models will also be re-evaluated prior to major updates that could materially enhance model capabilities.

Please see the “Evaluating Frontier Models for Critical Capabilities” section in our [Frontier Model Safety Framework](#).

---

**c. Describe how your organization conducts testing (e.g., red-teaming) to evaluate the model’s/system’s fitness for moving beyond the development stage?**

Amazon conducts extensive testing to evaluate frontier model fitness across the entire lifecycle of the model. As appropriate, we use a range of model evaluation approaches, including automated benchmarks, expert red teaming, and uplift studies.

We use multiple datasets and multiple human workforces to evaluate the performance of the Amazon Nova models. Our development testing involves automated benchmarking against publicly available datasets, automated benchmarking against proprietary datasets, benchmarking against proxies for anticipated customer use cases, human evaluation of completions against proprietary datasets, automated red teaming, manual red teaming, and more. Our development process examines Amazon Nova model performance using all these tests, takes steps to improve the model and/or the suite of evaluation datasets, and then iterates. In this [service card](#), we provide an overview of our methodology.

Amazon uses manual, automated and third parties to conduct red teaming against frontier models in areas such as safety, security, privacy, fairness, and veracity. We also work with specialized firms and academics to red-team frontier models for specialized areas such as Chemical, Biological, Radiological and Nuclear (CBRN) capabilities.

Please see the “Evaluating Frontier Models for Critical Capabilities” section in our [Frontier Model Safety Framework](#) and [Evaluating the critical risks of Amazon’s Nova Premier under the Frontier Model Safety Framework](#).

d. Does your organization use incident reports, including reports shared by other organizations, to help identify risks?

Yes

---

e. Are quantitative and/or qualitative risk evaluation metrics used and if yes, with what caveats? Does your organization make vulnerability and incident reporting mechanisms accessible to a diverse set of stakeholders? Does your organization have incentive programs for the responsible disclosure of risks, incidents and vulnerabilities?

Yes. *For Quantitative and Qualitative Risk Evaluations.* Please see our responses to questions 1(a) through 1(c).

*Reporting Mechanisms employed:*

o Amazon's automated threat intelligence and defense systems detect and mitigate millions of threats each day. These systems are backed by human experts for threat intelligence, security operations, and security response. Threat sharing with other providers and government agencies provides collective defense and response.

o The Amazon Cyber Threat Intelligence team continually monitors and tracks dozens of advance threat actor groups, observing their tactics, techniques, and procedures, and when appropriate, taking part in coordinated take-downs of their infrastructure. In addition, the AWS Trust & Safety and Fraud teams detect abusive and fraudulent behavior on the AWS cloud using automated and human monitoring, as well as external reporting mechanisms, and block or evict bad actors as needed.

o Amazon is continuing to invest in external security research, including bug bounty programs, academic research investments, and coordinated vulnerability disclosure programs that encourage and reward security experts to partner with us in research and development.

Please see "Appendix A: Amazon's Foundational Security Practices" in our [Frontier Model Safety Framework](#).

---

**f. Is external independent expertise leveraged for the identification, assessment, and evaluation of risks and if yes, how? Does your organization have mechanisms to receive reports of risks, incidents or vulnerabilities by third parties?**

Yes. Please see our responses to questions 1(a) through 1(e). In general, Amazon leverages external independent expertise where appropriate including proactive engagement across academic, industry, and government partnerships like the following:

- **Collaboration on threat modeling and updated Critical Capability Thresholds:** Amazon is committed to partnering with governments, domain experts, and industry peers to continuously improve our awareness of the threat environment and maintain Critical Capability Thresholds and evaluation processes that account for evolving and emerging threats.
- **Information sharing and best practices development:** Engagement in fora that bring together companies developing frontier models (e.g. Frontier Model Forum (FMF) and Partnership on AI) and organized by government agencies (e.g. National Institute of Standards and Technologies (NIST)). These fora serve as an opportunity to share findings and to adopt recommendations from other leading companies.
- **Fostering academic research for development of cutting-edge alignment techniques:** Through initiatives such as the [Amazon Scholars](#), [Amazon Research Awards](#) and Amazon Research centers (e.g. [USC + Amazon Center on Secure & Trusted Machine Learning](#), [Amazon/ MIT Science Hub](#)), we work with leading academic partners conducting research on frontier AI risks and novel risk mitigation approaches. Additionally, we advance our own research and publish findings in safety conferences, while borrowing learnings presented by other academic institutions at similar venues.
- **Investments in advanced AI safety R&D:** At Amazon, we accelerate our work in AI safety through initiatives such as our [Amazon AGI SF Lab](#) and the [Trusted AI Challenge](#). These channels enable us to leverage the work of subject matter experts and discover promising approaches towards aligning our frontier models.
- **Learning from our red teaming network:** We continue to build our strong network of internal and external red teamers including red teamers with deep subject matter expertise in risks related to critical capabilities. These experts are critical in surfacing early insights into emerging critical capabilities and help us identify and implement appropriate mitigations.

Through these partnerships and reporting mechanisms, we continuously enhance our ability to identify, assess, and address risks while contributing to the broader development of AI safety practices.

Please see our [Frontier Model Safety Framework](#) for more information.

**g. Does your organization contribute to the development of and/or use international technical standards or best practices for the identification, assessment, and evaluation of risks?**

Yes. Please see our responses to questions 1(e) and 1(f).

Amazon actively contributes to the development of international technical standards through active engagement in ISO/IEC's AI standards body (ISO/IEC JTC 1 SC 42) and AI-related projects in ISO/IEC's cybersecurity and privacy standards body (ISO/IEC JTC 1 SC 27). We also engage actively with other standards organizations like CEN/CENELEC, Coalition for Content Provenance and Authenticity (C2PA), the Institute of Electrical and Electronics Engineers (IEEE), and Internet Engineering Task Force (IETF). We recognize that effective standards establish common expectations about AI and responsible AI implementation. Standard implementations support our customers and both our upstream and downstream suppliers. We support AI safety and risk assessment and mitigation by engaging with/in organizations like NIST including the Artificial Intelligence Safety Institute Consortium (AISIC), Center for AI Standards and Innovation (CAISI), Thorn, and FMF.

---

**h. How does your organization collaborate with relevant stakeholders across sectors to assess and adopt risk mitigation measures to address risks, in particular systemic risks?**

Amazon actively engages with a diverse network of stakeholders, including academic institutions, industry partners, and government agencies. Amazon participates in cross-industry forums, supports academic research initiatives, and works with independent evaluators.

Please see our responses to questions 1(e) through 1(g).

---

## Section 2 - Risk management and information security

### a. What steps does your organization take to address risks and vulnerabilities across the AI lifecycle?

Amazon implements a comprehensive approach to risk mitigation that spans the entire AI lifecycle, from development through deployment and ongoing monitoring. Our framework includes technical safeguards, policy controls, and continuous evaluation processes designed to identify and address vulnerabilities while promoting safe and responsible AI development.

Please see the “Risk Mitigations: Safety and Security Measures for Frontier Models with Critical Capabilities” section in our [Frontier Model Safety Framework](#).

---

### b. How do testing measures inform actions to address identified risks?

Amazon's comprehensive testing measures, including automated benchmarks, expert red-teaming, and human-centric assessments, provide actionable insights that directly inform the implementation of specific safeguards and controls tailored to address identified risks.

Please see the “Risk Mitigations: Safety and Security Measures for Frontier Models with Critical Capabilities” section in our [Frontier Model Safety Framework](#).

---

### c. When does testing take place in secure environments, if at all, and if it does, how?

Testing of Amazon's frontier models takes place in secure compute environments designed to maintain data protection and operational security. Access to these environments is controlled through a comprehensive security framework.

Please see the “Appendix A: Amazon’s Foundational Security Practices” section in our [Frontier Model Safety Framework](#).

---

**d. How does your organization promote data quality and mitigate risks of harmful bias, including in training and data collection processes?**

Amazon implements a comprehensive approach to data quality and bias mitigation for its frontier models that spans pre-training, evaluation, and deployment.

For the Amazon Nova Micro, Lite, Pro, and Premier, the runtime service architecture works as follows: 1/ the model receives a user prompt via the API or Console; 2/ the model filters the prompt to comply with safety, fairness and other design goals; 3/ the model may augment the filtered prompt to support user-requested features, for example, knowledge-base retrieval; 4/ the model generates a completion; 5/ the model filters the completion for safety and other concerns; 6/ the model returns the final completion.

We say that an Amazon Nova model exhibits a particular "behavior" when it generates the same kind of completions for the same kinds of prompts with a given configuration (for example, temperature). For a given model architecture, the control levers that we have over the behaviors are primarily a/ the training data corpus, and b/ the filters we apply to pre-process prompts and post-process completions. Our development process exercises these control levers as follows: 1/ we pre-train the FM using curated data from a variety of sources, including licensed and proprietary data, open source datasets, and publicly available data where appropriate; 2/ we adjust model weights via supervised fine tuning (SFT) and reinforcement learning with human feedback (RLHF) to increase the alignment between the Amazon Nova model and our design goals; and 3/ we tune safety filters (such as privacy-protecting and profanity-blocking filters) to block or evade potentially harmful prompts and responses to further increase alignment with our design goals.

---

**e. How does your organization protect intellectual property, including copyright-protected content?**

We develop our products to respect privacy rights and support the protection of rights holders and content creators, as well as established legal frameworks that encourage/support/facilitate the development of innovative and beneficial services. For example, AWS offers uncapped intellectual property (IP) indemnity coverage for outputs of generally available Amazon Nova models (see Section 50.10 of the [AWS Service Terms](#)). This means that customers are protected from third-party claims alleging IP infringement or misappropriation (including copyright claims) by the outputs generated by these Amazon Nova models. In addition, our standard IP indemnity for use of the Services protects customers from third-party claims alleging IP infringement (including copyright claims) by the Services (including Amazon Nova models) and the data used to train them.

**f. How does your organization protect privacy? How does your organization guard against systems divulging confidential or sensitive data?**

Amazon Nova models are designed to avoid completing prompts that could be construed as requesting private information (PII). These models are available through Amazon Bedrock, a managed service that implements multiple layers of privacy protection. Amazon Bedrock does not store or review customer prompts or customer prompt completions, and prompts and completions are never shared between customers or with Amazon Bedrock partners. Additionally, AWS does not use inputs or completions generated through the Amazon Bedrock service to train Amazon Bedrock models, including Amazon Nova models. For more information, see Section 50.3 of the [AWS Service Terms](#) and the [AWS Data Privacy FAQs](#). For service-specific privacy information, see Security in the [Amazon Bedrock FAQs](#).

---

**g. How does your organization implement AI-specific information security practices pertaining to operational and cyber/physical security?**

- i. How does your organization assess cybersecurity risks and implement policies to enhance the cybersecurity of advanced AI systems?
- ii. How does your organization protect against security risks the most valuable IP and trade secrets, for example by limiting access to proprietary and unreleased model weights? What measures are in place to ensure the storage of and work with model weights, algorithms, servers, datasets, or other relevant elements are managed in an appropriately secure environment, with limited access controls in place?
- iii. What is your organization’s vulnerability management process? Does your organization take actions to address identified risks and vulnerabilities, including in collaboration with other stakeholders?
- iv. How often are security measures reviewed?
- v. Does your organization have an insider threat detection program?

Amazon’s approach to AI security is built on a firm, well-tested foundation of enterprise security controls for Amazon as a whole, and the unique, industry-leading security capabilities of the AWS cloud environment.

For detailed information about specific security implementations and protections, please see “Risk Mitigations: Safety and Security Measures for Frontier Models with Critical Capabilities” and Appendix A: Amazon’s Foundational Security Practices section in our [Frontier Model Safety Framework](#).

## h. How does your organization address vulnerabilities, incidents, emerging risks?

Amazon takes a lifecycle-based approach to managing the risks of AI, including potential risks that can arise from misuse of our advanced AI models. Our strategy is grounded in pre-deployment evaluation and post-deployment monitoring, with particular focus on three critical risk domains: Chemical, Biological, Radiological & Nuclear (CBRN), Offensive Cyber Operations, and Automated AI R&D.

Our pre-deployment evaluations combine automated benchmarks with human-centric assessments, including expert red-teaming and multi-agent stress tests. For example, our recent Amazon Nova Premier evaluation incorporated both reproducible automated testing and independent verification from specialized firms like Nemesys Insights for CBRN assessments and METR for automated AI R&D evaluations. These comprehensive assessments help our models remain within defined safety thresholds before release.

Post-deployment, we maintain continuous monitoring through automated detection systems and dynamic content filters that protect against misuse while preserving model functionality. For details about our security and privacy protections, please see our previous responses and the "Security" section of our [AWS AI Service Cards](#) for Amazon Nova models.

For more detailed information about our approach to specific risk domains and evaluation methodologies, please see our [Frontier Model Safety Framework](#) and [Amazon Nova Premier Evaluation report](#).

---

## Section 3 - Transparency reporting on advanced AI systems

a. Does your organization publish clear and understandable reports and/or technical documentation related to the capabilities, limitations, and domains of appropriate and inappropriate use of advanced AI systems?

- i. How often are such reports usually updated?
- ii. How are new significant releases reflected in such reports?
- iii. Which of the following information is included in your organization's publicly available documentation: details and results of the evaluations conducted for potential safety, security, and societal risks including risks to the enjoyment of human rights; assessments of the model's or system's effects and risks to safety and society (such as those related to harmful bias, discrimination, threats to protection of privacy or personal data, fairness); results of red-teaming or other testing conducted to evaluate the model's/system's fitness for moving beyond the development stage; capacities of a model/system and significant limitations in performance with implications for appropriate use domains; other technical documentation and instructions for use if relevant.

Yes. Amazon provides comprehensive documentation of our frontier models through multiple channels. For our Amazon Nova family of models, we publish detailed [AWS AI Service Cards](#), technical reports, and documentation that outline capabilities, limitations, and usage guidelines. Our recent Amazon Nova Premier evaluation report demonstrates this commitment to transparency, providing detailed assessments across critical capability domains including CBRN, Offensive Cyber Operations, and Automated AI R&D, safety evaluations and performance benchmarks.

For complete documentation, please see:

- [AWS AI Service Cards](#)
- [Amazon Nova Models Technical report and model card](#)
- [Amazon Nova Premier Evaluation report](#)

For more information on how Amazon Nova was built responsibly, please see [Amazon Nova and our commitment to responsible AI](#).

**b. How does your organization share information with a diverse set of stakeholders (other organizations, governments, civil society and academia, etc.) regarding the outcome of evaluations of risks and impacts related to an advanced AI system?**

We engage with a diverse set of stakeholders through multiple channels, including public documentation, academic partnerships, industry collaborations, government engagements, and independent evaluations. This approach allows us to contribute to and learn from the broader AI community while maintaining transparency about our models' capabilities, limitations, and safety assessments.

For detailed information about our sharing practices and stakeholder engagements please see our:

- [AWS AI Service Cards](#)
- [Amazon Nova Models Technical report and model card](#)
- [Amazon Nova Premier Technical report and model card](#)
- [Amazon Nova Sonic Technical report and model card](#)

[Amazon Nova Premier evaluation report](#) for an example of our comprehensive risk assessment and sharing practices.

---

**c. Does your organization disclose privacy policies addressing the use of personal data, user prompts, and/or the outputs of advanced AI systems?**

Yes. For complete documentation, please see:

- Section 50.3 of the [AWS Service Terms](#)
- [Amazon Bedrock FAQs](#)
- [Privacy Notice](#)
- [Data Privacy FAQ](#)
- [Amazon Model Training & Privacy](#)

---

**d. Does your organization provide information about the sources of data used for the training of advanced AI systems, as appropriate, including information related to the sourcing of data annotation and enrichment?**

Yes. Through our [AWS AI Service Cards](#) and technical documentation, we provide transparency about our data practices.

**e. Does your organization demonstrate transparency related to advanced AI systems through any other methods?**

Amazon demonstrates transparency through comprehensive technical documentation, practical safeguards in our frontier model development, and active engagement with the AI community. We publish detailed [AWS AI Service Cards](#) for our Amazon Nova Models, accompanied by technical reports and model cards that are updated with each new release. These documents provide clear information about model capabilities, limitations, and appropriate use cases.

Our transparency extends to practical safeguards in our frontier models. For example, our Amazon Nova Models (Reel, Canvas, and Sonic) incorporate robust content provenance mechanisms, including invisible watermarks that help identify AI-generated content. These features are complemented by detection solutions that allow users to verify whether content was generated by Amazon Nova models.

Our [Responsible AI](#) website offers transparency into our responsible AI strategy including our core dimensions of responsible AI, best practices, research and customer examples.

---

## **Section 4 - Organizational governance, incident management and transparency**

**a. How has AI risk management been embedded in your organization governance framework? When and under what circumstances are policies updated?**

Risk management is embedded throughout our organizational structure, anchored by our dedicated Responsible AI Governance team. This central team works in coordination with Responsible AI teams and other stakeholders across the company to drive consistent policy review, updates and compliance based on technological advances, emerging risks, and evolving best practices. Responsible AI focused teams across the company develop and communicate builder requirements and, as implementation continues to advance, assess the progress of builder teams against requirements and Responsible AI dimensions including fairness, explainability, privacy and security, safety, controllability, veracity and robustness, governance, and transparency.

---

**b. Are relevant staff trained on your organization’s governance policies and risk management practices? If so, how?**

Yes. Amazon provides numerous opportunities to develop AI literacy among our workforce and AI skill development for relevant staff. We have learning and development platforms to track employee training and ongoing education across various topics, including security awareness and inclusivity, which impact the quality of frontier models. For AI disciplines specifically, including responsible AI, governance and security, we provide: (a) custom training with service teams, (b) courses from AWS Machine Learning University (MLU), and (c) practical training via collaborations with scientists, Subject Matter Experts (SMEs) and training partners. Some of this internal training is also made available externally, including for customers. For example, [AWS MLU](#) provides a curriculum that spans AI foundational, technical, governance and risk management skillsets.

---

**c. Does your organization communicate its risk management policies and practices with users and/or the public? If so, how?**

Transparency is a core dimension of our approach to Responsible AI, shaping how we communicate our risk management policies and practices to users and the public. We believe effective risk management requires clear, accessible communication that enables stakeholders to make informed decisions about AI system deployment and use.

Amazon demonstrates its commitment to transparency through various channels, including:

- **Public Documentation:** We publish [AWS AI Service Cards](#), which include detailed information about our models' capabilities, limitations, and risk management practices. These are complemented by technical reports and model cards for our Amazon Nova Models.
- **User Guides:** Comprehensive [user guides](#) are available, providing instructions on responsible use of our frontier models and detailing built-in safeguards.
- **Policy Disclosures:** Our [AWS Service Terms](#), [Data Privacy FAQs](#), and service-specific FAQs (like the [Amazon Bedrock FAQs](#)) provide clear information about our data handling, privacy, and security practices.
- **Responsible AI web site:** Our [Responsible AI](#) website serves as a central hub for sharing our approach to responsible AI, including our risk management framework and customer real-world examples of responsible deployment.
- **Public Engagement:** We participate in industry forums and academic collaborations, sharing our approaches to AI safety and risk management.
- **Customer Communication:** For our enterprise customers, we provide more detailed information about our risk management practices through direct communications and support channels.

**d. Are steps taken to address reported incidents documented and maintained internally? If so, how?**

Yes. Amazon maintains processes and procedures for addressing reported incidents through established mechanisms. These procedures ensure consistent handling of incidents and enable continuous improvement of our response capabilities. This includes procedures for triage and escalation, ensuring systematic review, validation, and resolution of reported issues.

These processes are supported by cross-functional collaboration, where abuse insights, intelligence, and emerging risks are shared across teams. This allows us to track patterns, identify emerging risks, and continuously improve our prevention and response strategies while maintaining appropriate security controls.

---

**e. How does your organization share relevant information about vulnerabilities, incidents, emerging risks, and misuse with others?**

Amazon employs a comprehensive approach to sharing information about vulnerabilities, incidents, emerging risks, and potential misuse through reporting and notification mechanisms and cross-industry collaboration. Our strategy balances the need for transparency with the imperative to maintain robust security measures. For more detailed information, please see:

- The "Security" section in our AWS AI Service Cards for [Amazon Nova Micro, Lite, Pro, and Premier](#) and [Amazon Nova Sonic](#).
- Amazon's [Frontier Model Safety Framework](#).

---

**f. Does your organization share information, as appropriate, with relevant other stakeholders regarding advanced AI system incidents? If so, how? Does your organization share and report incident-related information publicly?**

Amazon takes a structured approach to sharing incident-related information about our frontier models, balancing the need for transparency with security considerations. Through multiple channels and partnerships, we engage with relevant stakeholders to share insights about incidents, emerging risks, and mitigation strategies, contributing to the broader advancement of AI safety practices while maintaining appropriate security protocols.

Please see Information sharing and best practices development of the [Amazon Frontier Model Safety Framework](#).

**g. How does your organization share research and best practices on addressing or managing risk?**

Amazon shares research and best practices through multiple channels.

- We actively contribute to industry knowledge through engagement in fora that bring together companies developing frontier models including the FMF and Partnership on AI. Our participation in platforms organized by government agencies like NIST creates opportunities to share findings and adopt recommendations from other leading companies.
- AWS MLU, in collaboration with the Worldwide Specialist Organizations AI/ML Technical Field Community and Professional Services, develop and share technical training resources and best practices.
- We publish our own research findings including technical reports for [Amazon Nova Premier](#) and [Amazon Nova Sonic](#) as well as the detailed [evaluations](#) under Amazon's Frontier Model Safety Framework.
- Additionally, through partnerships with academic institutions, we advance research through initiatives like the Amazon Nova Trusted AI Challenge and [Amazon Research Awards](#), and through dedicated research centers such as the [USC + Amazon Center on Secure & Trusted Machine Learning](#) and the [Amazon/MIT Science Hub](#).

---

**h. Does your organization use international technical standards or best practices for AI risk management and governance policies?**

Yes. We use international technical standards or best practices for AI risk management as helpful to support our internal governance practices.

---

## **Section 5 - Content authentication & provenance mechanisms**

**a. What mechanisms, if any, does your organization put in place to allow users, where possible and appropriate, to know when they are interacting with an advanced AI system developed by your organization?**

Amazon's Nova models implement transparency mechanisms that enable stakeholders to make informed choices about their engagement with an AI system.

For example, Amazon Nova models integrate content provenance, watermarking, and detection mechanisms to provide transparency to our customers.

b. Does your organization use content provenance detection, labeling or watermarking mechanisms that enable users to identify content generated by advanced AI systems? If yes, how? Does your organization use international technical standards or best practices when developing or implementing content provenance?

On September 12, 2024, [Amazon announced it was joining the C2PA as a Steering Committee member](#). Amazon attaches Content Credentials to the visual assets produced using the company's generative AI models Nova Canvas and Titan Image Generator v1 and v2, which allows enterprise customers to create and edit images in a range of visual styles. Amazon is also working to incorporate Content Credentials in AWS Elemental MediaConvert, a file-based video processing service that transcodes content for broadcast and multi-screen delivery at scale. This implementation will allow its customers - such as news organizations and sports broadcasters - to communicate the provenance of the media they are sharing, allowing their distribution partners and news aggregators to verify the content's authenticity before posting it on their platforms.

[Amazon Nova Reel](#), [Amazon Nova Canvas](#), and [Amazon Nova Sonic](#) all have content provenance watermarking for the safe and responsible use of AI.

- Amazon Nova Reel applies an invisible watermark to all videos it generates, helping identify AI-generated videos to promote the safe, secure, and transparent development of AI technology and helping reduce the spread of disinformation.
  - Amazon Nova Canvas applies an invisible watermark to all images it generates, helping identify AI-generated images to promote the safe, secure, and transparent development of AI technology and helping reduce the spread of disinformation.
  - Amazon Nova Sonic embeds a robust, nearly imperceptible watermark into all generated speech outputs. This watermark enables reliable attribution of generated audio to Amazon Nova Sonic.
-

# Section 6 - Research & investment to advance AI safety & mitigate societal risks

a. How does your organization advance research and investment related to the following: security, safety, bias and disinformation, fairness, explainability and interpretability, transparency, robustness, and/or trustworthiness of advanced AI systems?

Amazon advances research and investment in AI safety through multiple channels and partnerships. In collaboration with the National Science Foundation (NSF), we jointly funded a \$20 million, three-year program focused on [Fairness in AI](#). This initiative has supported 70 faculty members across more than 40 institutions, resulting in 283 publications that advance our understanding of AI fairness and safety. We've committed \$5 million in compute credits to the [NIST U.S. Artificial Intelligence Safety Institute Consortium](#) to enable the development of tools and methodologies to evaluate the safety of foundation models, and £5 million to the international coalition led by United Kingdom AI Security Institute AI Alignment Project tackling the challenge of AI systems behaving predictably and alignment with human values.

Our commitment to academic research is further demonstrated through the [Amazon Research Awards](#) program, which offers unrestricted funds and AWS Promotional Credits to support research at academic institutions and non-profit organizations. This includes a dedicated Fairness in AI track with approximately \$1 million in funding. We also award annual cybersecurity research grants that encompass generative AI security research.

Through our [University Hubs Program](#), we foster collaboration between academia and industry, making research findings publicly accessible and presenting them at public events. Currently, we maintain fourteen Hubs at institutions including Columbia University, University of Texas Austin, University of South Carolina and the Max Planck Institute in Germany. Since establishing our first Hub at Columbia in 2021, the program has funded over 120 Sponsored Research Projects and 65 PhD Fellowships.

---

b. How does your organization collaborate on and invest in research to advance the state of content authentication and provenance?

Amazon collaborates with external organizations to advance research in content authentication and provenance as evidenced by our involvement in organizations like the Coalition for Content Provenance and Authenticity (C2PA) through which we help develop and implement technical standards for content authentication and provenance.

---

**c. Does your organization participate in projects, collaborations, and investments in research that support the advancement of AI safety, security, and trustworthiness, as well as risk evaluation and mitigation tools?**

Amazon demonstrates an ongoing commitment to advancing AI safety through both internal programs and external collaborations. Through the Amazon Research Awards and dedicated research centers like the USC + Amazon Center on Secure & Trusted Machine Learning and the Amazon/MIT Science Hub, we fund and conduct foundational research on frontier AI risks and novel mitigation approaches. Our collaboration with the National Science Foundation on the Fairness in AI program represents a \$20 million investment that has supported 70 faculty members across more than 40 institutions, resulting in 283 publications advancing AI fairness and safety.

Our recent work with Amazon Nova Premier exemplifies our structured approach to safety research through several key initiatives. Please see [here](#) for more information.

Through our [University Hubs Program](#), we've funded over 120 Sponsored Research Projects and 65 PhD Fellowships since 2021. Our [Amazon AGI Lab](#) and [Trusted AI Challenge](#) accelerate work in AI safety by leveraging subject matter experts to develop innovative approaches for aligning frontier models.

These research initiatives are complemented by our active participation in industry forums and standards organizations, where we contribute to the development of risk evaluation and mitigation tools while sharing insights with the broader AI safety community. This approach advances the field of AI safety while maintaining rigorous standards for model deployment, as demonstrated by our commitments at the 2025 Paris AI Safety Summit.

---

#### d. What research or investment is your organization pursuing to minimize socio-economic and/or environmental risks from AI?

Amazon pursues research and investment initiatives aimed at minimizing both socio-economic and environmental risks from AI development and deployment.

Amazon is taking a multifaceted approach to balancing AI advancement with environmental and social responsibility. At the infrastructure level, AWS introduced new data center components offering 12% more compute power with improved efficiency, along with direct-to-chip liquid cooling solution for high-density AI compute chips that reduces mechanical energy consumption by up to 46% during peak cooling—without increasing water usage. These improvements helped AWS achieve a global Power Usage Effectiveness (PUE) of 1.15, compared to the industry average of 1.25.

We're investing in projects to reduce the water footprint of our facilities and expand water availability in the communities where we operate. In 2022, we announced our commitment to being [water positive by 2030](#) returning more water to communities and the environment than we use in AWS data center operations by 1) increasing the use of sustainable water sources, 2) improving water use efficiency across our operations, 3) reusing water as much as possible, and 4) supporting water replenishment projects for communities and the environment around the world. In partnership with the Water Environment Federation, The Water Center at the University of Pennsylvania, and Leading Utilities of the World we established the [Water-AI Nexus Center of Excellence](#). This first-of-its kind initiative leads and convenes research at the intersection of water and AI, advances AI-powered community water solutions and empowers, the next generation of diverse water leaders in the AI era.

To address the increased energy demands of AI, we focus on optimizing efficiency while scaling carbon-free energy. AI is being used to optimize sizing recommendations to reduce returns, identifying energy inefficiencies, detecting water leaks, and avoiding packaging; solving countless environmental challenges while simultaneously improving service quality. Please see our [Amazon Sustainability Report](#).

Through our [Clean Energy](#) and [Sustainability](#) accelerators, we invest in start-ups, scale-ups, and SMBs driving sustainable innovation. These programs facilitate co-innovation between leading Energy, Utilities & Industrial corporations and clean energy & climate tech innovators. We help entrepreneurs develop their skills and scale their businesses to maximize climate impact, while developing breakthrough technologies to support net-zero goals. In partnership with UNESCO and the International Centre on Artificial Intelligence (IRCAI), our [Compute for Climate Fellowship](#) funds proof of concepts climate solutions that leverage advanced cloud computing and generative AI addressing the climate crisis.

Additionally, the [AWS Imagine Grant](#) program exemplifies our commitment to addressing societal challenges through AI. The program's Pathfinder - Generative AI awards support innovative projects like The Nature Conservancy's development of the first-ever generative AI suite of tools for nature-based science, and RAINN's implementation of AI-powered crisis hotlines for victims of abuse.

---

## Section 7 - Advancing human and global interests

a. What research or investment is your organization pursuing to maximize socio-economic and environmental benefits from AI? Please provide examples.

Amazon leverages our extensive experience in AI development to deliver socio-economic and environmental benefits by lowering barriers to innovation, funding public-interest projects, and accelerating climate and equity solutions. Our efforts include major funding commitments to socioeconomic initiatives and accelerator programs supporting climate-tech and clean energy startups.

Specific examples of our work include:

- AI-Driven Sustainability: We've developed FlowMS, an AI-powered solution that enhances [building energy and water efficiency](#), accelerating business decarbonization and improving operational efficiency. This technology has been implemented across our global operations, contributing to our achievement of a 1.15 Power Usage Effectiveness (PUE) rating, significantly better than the industry average of 1.25.
- Public Sector Innovation: [The AWS Public Sector Generative AI Impact Initiative](#) funds projects that help agencies modernize constituent services. This includes improvements in case management, language accessibility, and information delivery, directly supporting the digital transformation of public services.
- Health Equity: Through the [The AWS Health Equity Initiative](#), we support AI projects that reduce disparities in care, expand telehealth, improve diagnostics, and increase access for underserved communities. This program has funded numerous projects across multiple countries, improving healthcare access for millions.
- Disaster Response: [AWS Disaster Response](#) partners with emergency responders, leveraging AI and geospatial analysis for rapid damage assessment, flood mapping, and resource allocation. This technology has been deployed in multiple natural disaster scenarios, significantly enhancing response capabilities.
- AI Education: Our [Amazon's AI Ready](#) initiative provides free AI skills training to 2 million people globally through 2025 and complements the [AWS Education Equity Initiative](#), which empowers organizations to create digital learning solutions for underserved learners worldwide. Amazon has also joined education, technology, and government leaders investing in AI education for K-12 students and educators as part of the [White House's Pledge to America's Youth](#).
- Sustainable Development: We support the [AI Hub for Sustainable Development](#), designed to accelerate skill development and AI adoption to power industrial growth in Africa. This initiative aims to bridge the global AI skills gap and foster economic development in emerging markets.
- Climate Innovation: In partnership with UNESCO and the International Research Centre on Artificial Intelligence (IRCAI), our Compute for Climate Fellowship funds proof-of-concept climate solutions that leverage advanced cloud computing and generative AI to address the climate crisis.

**b. Does your organization support any digital literacy, education or training initiatives to improve user awareness and/or help people understand the nature, capabilities, limitations and impacts of advanced AI systems? Please provide examples.**

Yes. Amazon has invested in and launched several educational initiatives to drive AI literacy and understanding globally. Through our "[Amazon AI Ready](#)" initiative, we've committed to providing free AI skills training to 2 million people globally by 2025. This includes launching new free training courses about safe and responsible AI use through our digital learning centers, with specific courses like "Introduction to Responsible AI" for new-to-cloud learners on AWS Educate and more advanced offerings such as "Responsible AI Practices" and "Security, Compliance, and Governance for AI Solutions" on AWS Skill Builder.

Our educational outreach extends across multiple programs designed to address different needs and skill levels:

- The [AWS Education Equity](#) Initiative empowers organizations to create digital learning solutions for underserved learners globally
- [AWS Skill Builder](#) provides online learning resources for all skill levels offering more than 100 AI-focused courses
- [AWS re/Start](#) we help unemployed and underemployed individuals launch new careers in cloud technology, with specific focus on cloud, AI and machine learning skills.
- [AWS AI and ML Scholarships](#) enable students to gain hands-on artificial intelligence and machine learning training
- [Amazon Future Engineer](#) helps students worldwide explore career pathways in computer science and emerging technology
- [AWS MLU](#) makes our internal AI training curriculum available to the public, sharing the same resources we use to train Amazon's own developers and scientists.
- [AWS AI Educator Enablement Program](#) equips community colleges and minority-serving institutions develop the skills to equip students with AI/ML entry-level skills.

These initiatives reflect our commitment not just to advancing AI technology, but to ensuring broad understanding of both its capabilities and limitations across diverse communities.

**c. Does your organization prioritize AI projects for responsible stewardship of trustworthy and human-centric AI in support of the UN Sustainable Development Goals? Please provide examples.**

Yes. Amazon works with customers globally to harness our technology in addressing challenges aligned with the UN's Sustainable Development Goals (SDGs). We demonstrate this commitment through concrete initiatives and programs. For example, in 2022, we launched the [Amazon Sustainability Data Initiative](#) (ASDI) in collaboration with UNESCO's International Research Centre in Artificial Intelligence, challenging participants to develop sustainability solutions using ASDI data on AWS Cloud services to support specific SDGs.

---

**d. Does your organization collaborate with civil society and community groups to identify and develop AI solutions in support of the UN Sustainable Development Goals and to address the world's greatest challenges? Please provide examples.**

Yes. Amazon collaborates with various organizations to develop AI solutions addressing global challenges. We're committed to unlocking AI's potential for social good by reducing barriers that social enterprises often face due to limited resources. For example, our support for the [AI for Changemakers Accelerator](#) program, led by Tech To The Rescue (TTTR) helps social impact organizations access necessary AI tools and expertise to scale their impact projects. [AWS Imagine Grants](#) for Nonprofits Pathfinder Generative AI provides unrestricted financial support, credits and implementation support for nonprofit driving their mission with Generative AI.

Our collaborations have produced notable results in supporting SDGs:

- Working with Jacaranda Health in Kenya to deploy an AI-enabled help desk that improves maternal health outcomes, now reaching 1.5 million mothers
- Supporting Growy, a Netherlands-based startup, in operating fully automated farms using AWS IoT Events to collect and analyze data for food security
- Partnering with organizations like Thorn and All Tech is Human to safely design generative AI services that reduce child exploitation risks

These partnerships demonstrate our commitment to leveraging AI technology to address critical global challenges through responsible development and deployment of AI solutions.