



CPSC Artificial Intelligence Strategy

Introduction

The U.S. Consumer Product Safety Commission (CPSC) is committed to responsibly accelerating the adoption of Artificial Intelligence (AI) to advance its mission of protecting the public from unreasonable risks of injury associated with consumer products. AI presents an unprecedented opportunity to transform how CPSC identifies emerging hazards, prioritizes risks, and delivers timely, data-driven actions that better protect consumers.

To realize this opportunity, CPSC is modernizing its data, analytics, and decision-making capabilities through cloud-based infrastructure, modern data architecture, advanced analytics, and responsible AI. These investments will enable the agency to integrate diverse information sources more efficiently, generate actionable insights more rapidly, and improve the speed, scale, and consistency of safety analyses. Supported by strong governance and human oversight, these capabilities will enhance CPSC's ability to identify emerging hazards earlier, inform regulatory and operational decisions, and respond more effectively to an evolving product safety landscape.

This AI Strategy¹ establishes CPSC's enterprise approach for expanding the use of AI to improve hazard identification, risk assessment, regulatory decision-making, operational efficiency, and scientific analysis while ensuring AI systems are implemented in a secure, transparent, and trustworthy manner. Developed in accordance with Executive Order 14179, *Removing Barriers to American Leadership in Artificial Intelligence*, and Office of Management and Budget (OMB) Memoranda M-25-21, *Accelerating Federal Use of AI through Innovation, Governance, and Public Trust*, and M-25-22, *Driving Efficient Acquisition of Artificial Intelligence in Government*, this strategy outlines the agency's priorities for AI innovation, governance, infrastructure modernization, workforce development, and risk management, while building upon CPSC's broader data modernization efforts to strengthen consumer product safety and enhance public trust.

CPSC AI Use Cases

CPSC is applying AI across mission and mission-support functions to improve product safety, strengthen operational efficiency, and enhance decision-making. AI use cases span the product safety lifecycle, including hazard identification, data processing, import surveillance, regulatory operations, and administrative functions. These efforts leverage machine learning, predictive analytics, generative AI, and intelligent automation to improve the speed, consistency, and scalability of agency operations while maintaining appropriate human oversight.

CPSC has established a strong foundation for enterprise AI adoption through the deployment of Microsoft Copilot and other approved AI capabilities. As an early federal adopter of Microsoft Copilot, the agency is already using generative AI to improve productivity across administrative, analytical, software development, and knowledge management activities. In parallel, CPSC has

¹ This strategy document was prepared by the CPSC staff. It has not been reviewed or approved by, and may not represent the views of, the Commission.



United States Consumer Product Safety Commission

developed mission-specific machine learning capabilities, including machine learning models for product classification and data imputation, demonstrating the agency's commitment to responsibly integrating AI into both mission-support and consumer product safety operations. Building upon this foundation, CPSC developed an enterprise AI Use Case Inventory to identify, prioritize, and govern future AI investments.

CPSC developed its AI Use Case Inventory through a collaborative, agency-wide effort involving all program offices. Offices were invited to identify operational challenges that lend themselves to an AI solution. Each proposed use case was documented using the Heilmeier Catechism to clearly articulate the problem being addressed, the proposed AI solution, expected benefits, technical feasibility, implementation considerations, and measures of success.

The resulting inventory was then evaluated and prioritized based on factors such as mission impact, strategic alignment, technical feasibility, data readiness, expected return on investment, implementation complexity, resource requirements, and potential risks. This structured approach enables CPSC to focus investments on high-value AI capabilities while maintaining a transparent and repeatable process for evaluating future AI use cases.

Representative AI use cases include:

- **Software Development Assistance** – AI coding assistants will be used to assist developers with code generation, testing, debugging, documentation, and modernization of legacy applications. This accelerates the delivery of secure, reliable digital services that support CPSC's mission. AI-generated code will be reviewed, tested, and approved prior to deployment.
- **Intelligent Data Ingestion and Integration** – Automate information classification tasks, including extracting, summarizing, and routing information from a variety of structured and unstructured data sources while maintaining current quality assurance and quality control checks. This will reduce manual processing efforts and enable faster identification and analysis of potential product safety hazards.
- **AI-Powered Knowledge Assistant** – Provide an AI-enabled query assistant that enables employees to search, analyze, and synthesize information from enterprise systems subject to established security and access controls. Query assistants will enhance workforce productivity and support timely, informed regulatory decision-making for compliance, enforcement, and risk reduction.
- **Injury Data Analysis** – Streamline analysis and summarization of injury data to improve the speed, consistency, and scalability of safety and injury surveillance in accordance with applicable data privacy and usage agreements. Automated analysis and summarization will enable earlier detection of emerging hazards and strengthen the scientific basis for product safety decisions.
- **Risk-Based Import Screening** – Improve the identification of high-risk imported consumer products and enhance inspector decision-making to prevent hazardous imported products from reaching consumers.
- **Sample Evaluation Assistant** – Streamline the preparation, summarization, evaluation of evidence collected by Compliance Officers to determine if a sample product is violative or



United States Consumer Product Safety Commission

presents a substantial product hazard. Assistants will reduce the amount of time that Compliance Officers spend preparing and analyzing cases.

CPSC AI Maturity Goals

CPSC has and will continue advancing its AI capabilities through a coordinated strategy that strengthens governance, workforce capabilities, data infrastructure, and the responsible adoption of AI. The agency's goal is to make AI a core operating capability by redesigning business processes and workflows so people and AI systems operate together in a secure, efficient, and mission-focused environment.

Achieving this vision requires embedding automation where it delivers measurable value, maintaining human oversight where judgment and accountability are essential, and building systems, data, and governance with AI in mind from the outset. This means rethinking how work is performed: embedding automation where it delivers clear value, reinforcing human oversight where judgment is essential, and ensuring that systems, data flows, and governance structures are built for AI from the beginning. CPSC has and will continue improving:

- **AI-enabling infrastructure** – cloud-native environments, secure Application Programming Interfaces (APIs), modern analytical platforms.
- **High-quality, centralized data** – a data lake with integrated, governed, normalized, domain-specific datasets with the requisite semantic layer that enable reliable AI and Machine Learning (ML).
- **An AI-ready workforce** – staff trained to use and oversee AI tools like Copilot, supported by technical experts in data engineering, data science, and ML.
- **Strong governance and human oversight** – standards, risk controls, validation processes, and continuous monitoring for trustworthy AI use.

AI-Enabling Infrastructure

CPSC is operating on a cloud-native infrastructure that enables the agency to become fully AI-native. The agency will continue investing in enterprise infrastructure that enables the responsible development, deployment, operation, and governance of AI systems. To support enterprise AI capabilities, CPSC will continue to modernize the following infrastructure components:

- **Cloud Computing Environment:** Expand cloud-native computing resources to support AI model development, training, deployment, and operational workloads. Cloud services provide the scalability, reliability, security and computational resources necessary for modern AI applications.
- **Enterprise Data Platform:** Continue expanding the CPSC Data Lake and related data integration capabilities to provide trusted, governed, and accessible data for AI development and analytics. Enterprise data management capabilities will improve data quality, interoperability, lineage, and availability across mission systems.



- **Open-Source Analytics Environment:** Continue the transition toward standardized, cloud-supported analytical environments to support data science, machine learning, statistical analysis, and generative AI development. Modern open-source tools will provide greater flexibility, reproducibility, and long-term sustainability while reducing dependence on legacy analytical platforms.
- **Enterprise AI Platforms and Services:** Leverage approved enterprise AI services, including commercial and government-provided AI platforms, to support generative AI, document analysis, natural language processing, software development, and other mission applications. AI services will be evaluated and deployed in accordance with agency governance, cybersecurity, privacy, and acquisition requirements.
- **Application Programming Interfaces and System Integration:** Expand the use of secure APIs and standardized integration services that enable AI applications to interact with enterprise systems, external data sources, and mission applications while maintaining appropriate security and governance controls.
- **Secure Development and Operations:** Continue implementing modern software engineering, Machine Learning Operations (MLOps), Development, Security, and Operations (DevSecOps), continuous integration and continuous deployment (CI/CD), model monitoring, and cybersecurity practices that support the secure development, deployment, maintenance, and continuous improvement of AI-enabled systems.

Data

Data is foundational to CPSC's AI Strategy. CPSC will continue to modernize its enterprise data ecosystem to ensure AI systems are built on trusted, high-quality, and well-governed data. Through an enterprise Data Lake, standardized data architectures, automated ingestion pipelines, interoperable systems, and robust data governance, CPSC is creating an analytics-ready environment that supports scalable AI, data traceability, and faster data-driven decisions that strengthen consumer product safety.

As a foundational piece of CPSC's AI Strategy, CPSC will:

- Modernize data engineering and ingestion pipelines to automate the acquisition, validation, transformation, and integration of structured and unstructured data, creating scalable, AI-ready data pipelines for analytics and decision-making.
- Utilize a "Data Lake-first" approach that integrates data from internal and external sources.
- Normalize structured and unstructured data and curate domain-specific datasets.
- Enhance data governance and create a semantic layer by creating an inventory of enterprise data assets; developing consistent business definitions; creating and tagging metadata; documenting data provenance, lineage, and transformations to ensure discoverability, consistency, traceability, and reusability of trusted data for AI applications.
- Apply explainability and interpretability principles by maintaining documentation for AI models, training data, testing, validation, and performance monitoring to support human oversight and transparent decision-making.



United States Consumer Product Safety Commission

- Promote the secure sharing of data and AI assets across the agency and, where appropriate, with other federal agencies, researchers, and the public through established governance processes, data-sharing agreements, and public data resources.

AI Ready Workforce

CPSC is committed to building an AI-ready workforce that can confidently and responsibly integrate AI into every aspect of its mission. The agency has established a strong foundation through enterprise AI training, technical workforce development, and collaborative communities of practice, and will continue investing in the skills, tools, and culture needed to enable widespread AI adoption.

CPSC's vision is a workforce where employees at every level understand how to effectively leverage AI in their daily work, supported by secure, mission-focused AI assistants and modern analytical environments. By fostering AI literacy, expanding technical expertise, and encouraging responsible experimentation, CPSC is preparing its workforce to collaborate seamlessly with AI, improve decision-making, increase operational efficiency, and accelerate consumer product safety while maintaining the highest standards of scientific integrity, transparency, and accountability.

As a part of this operating model, CPSC has designated Microsoft Copilot as its official in-tenant generative AI platform, enabling staff to safely leverage AI for drafting, summarization, and analysis within the agency's secure Microsoft 365 environment. As AI capabilities continue to evolve, CPSC will continue to invest in a workforce with the knowledge, skills, and governance necessary to responsibly leverage AI at CPSC while maintaining appropriate human oversight and accountability:

- Continue to provide role-based AI training tailored to executives, managers, technical staff, acquisition personnel, and general users.
- Continue developing technical expertise in AI, machine learning, data science, cloud computing, and modern analytical tools through training, communities of practice, strategic hiring, and hands-on experience.
- Foster a culture of innovation and continuous learning by encouraging collaboration, knowledge sharing, experimentation, and the responsible adoption of emerging AI capabilities across the agency.
- Promote responsible AI use by providing guidance on human oversight, validation of AI-generated outputs, cybersecurity, privacy, records management, and ethical AI practices.
- Leverage government-wide partnerships, professional development opportunities, and internal knowledge-sharing networks to continuously strengthen the agency's AI capabilities and prepare the workforce for future advances in AI.

Research & Development

CPSC will continue researching and evaluating AI technologies that improve consumer product safety through enhanced hazard identification, data processing, and scientific analysis. Research and development efforts will focus on practical, mission-driven applications that can be validated and responsibly transitioned into operational use.



United States Consumer Product Safety Commission

The agency will continue expanding its existing machine learning models to automate the integration of additional data sources, including EHRs and other unstructured product safety data. CPSC will also conduct pilot projects to evaluate emerging AI technologies and collaborate with other Federal agencies, academia, and industry to leverage advances in AI that support the agency's mission.

Governance, Risk Management, and Resource Management

CPSC is committed to ensuring that AI systems are developed, acquired, deployed, and operated in a responsible, transparent, secure, and lawful manner. The agency's enterprise AI governance framework and policies will promote innovation while providing oversight, accountability, and public trust throughout the AI and ML lifecycle.

The CPSC Chief AI Officer (CAIO), in coordination with the AI Governance Board, provides enterprise oversight for AI strategy, governance, compliance, and risk management. The AI Governance Board reviews AI use cases, establishes agency-wide policies and standards, oversees the AI portfolio, and coordinates with agency stakeholders responsible for data governance, cybersecurity, privacy, acquisition, records management, and legal compliance.

Governance - CPSC will:

- Maintain an enterprise AI governance framework defining roles, responsibilities, and oversight across the AI lifecycle.
- Review AI use cases to ensure alignment with agency priorities, legal authorities, and responsible AI principles.
- Maintain the agency's AI Use Case Inventory and annually review and update the AI Strategy, policies, and governance processes to reflect evolving federal guidance, performance insights, and stakeholder input.
- Ensure transparency by publicly releasing the AI Strategy and Use Case Inventory and communicating significant updates as part of the annual review cycle.
- Align AI procurement and acquisition planning with federal requirements by evaluating opportunities to reuse existing capabilities, considering U.S.-developed AI solutions, incorporating appropriate technical and risk-management requirements into acquisitions, and engaging AI governance, acquisition, cybersecurity, privacy, legal, and IT officials throughout the acquisition lifecycle.

Risk Management – CPSC will:

- Apply an AI Risk Management Framework, such as the National Institute of Standards and Technology (NIST) AI Risk Management Framework, to guide the governance, development, acquisition, deployment, and monitoring of AI systems throughout their lifecycle.
- Apply a risk-based approach to classifying AI systems and implement safeguards appropriate to their level of risk. CPSC will determine High-Impact AI status using the federal criteria established in Office of Management and Budget guidance.



United States Consumer Product Safety Commission

- Require High-Impact AI systems to follow Risk Management policies:
 - Require High-Impact AI systems to implement a documented risk management process throughout the AI lifecycle, including risk identification, assessment, mitigation, acceptance, and continuous monitoring.
 - Conduct enhanced testing and evaluation of High-Impact AI systems prior to deployment and throughout operation to validate performance, reliability, security, and fitness for intended use.
 - Maintain documentation for High-Impact AI systems, including testing results, risk assessments, design decisions, performance metrics, and other records necessary to support transparency, oversight, and accountability.
 - Implement appropriate human oversight and governance to ensure that High-Impact AI systems remain subject to meaningful review, intervention, and operational controls throughout their lifecycle.
 - Develop contingency and response procedures for High-Impact AI systems that address model failures, performance degradation, security incidents, or other operational risks, including processes for corrective actions, rollback, or system suspension when appropriate.

Continuously monitor deployed AI systems for performance, security, reliability, and unintended impacts throughout their lifecycle.

Resource Management – CPSC will:

- Prioritize AI investments based on mission value, operational impact, risk, and available resources.
- Reuse enterprise AI capabilities and shared infrastructure where practical to improve efficiency and reduce duplication.
- Evaluate the AI portfolio periodically to identify opportunities for expansion, modernization, or retirement.

Through effective governance, risk management, and resource stewardship, CPSC will ensure that AI technologies advance the agency's mission while maintaining transparency, accountability, security, and public trust.

Brien Lorenze
Executive Director
Chief Artificial Intelligence Officer

August 25, 2026

Date