U.S. Department of Commerce (DOC)
Compliance Plan for OMB Memorandum M-25-21

September 2025

Prepared by Brian Epley, Chief Information Officer (CIO) and Chief AI Officer (CAIO)

Purpose

The Office of Management and Budget (OMB) Memorandum M-25-21, "Accelerating Federal Use of AI through Innovation, Governance, and Public Trust", directs each agency to submit to OMB and post publicly on the agency's website either a plan to achieve consistency with M-25-21, or a written determination that the agency does not use and does not anticipate using covered AI.

In accordance with this requirement, the Department of Commerce (Department or DOC) has developed its AI Compliance Plan, which outlines the Department's activities and compliance related to M-25-21's main goals of:

1. Driving Innovation
2. Improving AI Governance
3. Fostering Public Trust in Federal Use of AI

Driving AI Innovation

Removing Barriers to the Responsible Use of AI

DOC has identified select structural and operational barriers to the responsible adoption of AI and is undertaking targeted actions to mitigate or eliminate these barriers in alignment with OMB Memorandums M-25-21 and M-25-22.

Workforce Gaps:

DOC recognizes that a shortage of personnel with advanced AI competencies presents a barrier to responsible and effective AI adoption. To address this, the Department, in coordination with the Chief AI Learning Officer and the Office of the Chief Information Officer (OCIO), is expanding enterprise-wide training and professional development opportunities. This includes the development of a DOC-wide AI learning curriculum that covers foundational and applied topics such as machine learning, deep learning, natural language processing, and computer vision. These efforts are designed to upskill the current workforce, build long-term AI talent pipelines, and ensure compliance with Federal workforce readiness requirements outlined in M-25-21.

Infrastructure Limitations:

Modernization of IT infrastructure is essential to ensure scalable and secure AI development, testing, deployment, and monitoring. DOC is actively upgrading enterprise platforms and leveraging standardized governance processes to accelerate responsible prototyping and deployment cycles. These efforts include improving access to enterprise-approved AI tools, open-source libraries, and cloud-based environments to support experimentation and operationalization in a manner that upholds Federal cybersecurity and privacy requirements.

Policy and Procurement Constraints:

Legacy acquisition processes can inhibit timely and responsible access to emerging AI technologies. To mitigate this, DOC is aligning its procurement and acquisition practices with the guidance in M-25-22 to streamline contracting, ensure responsible vendor sourcing, and reduce administrative barriers while maintaining strict adherence to Federal procurement regulations. These measures will enhance DOC's ability to acquire, test and integrate AI solutions in a manner that is both agile and accountable.

Sharing and Reuse of AI Assets

DOC leverages a structured AI governance process to bring together AI leaders from across the agency to identify opportunities and promote and encourage the internal reuse of AI models and code, including Bureau AI Leads from across the agency and subject matter experts in related areas including acquisitions, budget and finance, cyber security, data, enterprise architecture, human resources, investigations, legal, privacy, security, technology and training.

The governance process utilizes a central repository of AI use cases to identify opportunities for reuse and encourage collaboration on those specific opportunities. In addition, DOC focuses on leveraging preferred, readily available solutions such as its Microsoft ecosystem to manage and share data. The DOC Chief Information Officer (CIO) oversees the coordination of AI assets and ensures alignment with open government principles.

To further enable and accelerate this activity DOC would benefit from additional AI subject matter expertise and project management resources.

AI Talent

The Department of Commerce is actively working on several initiatives to enhance AI talent and ensure our workforce is well-equipped to leverage AI technologies effectively that supports OMB M-25-21 and M-25-22. The Office of Chief Information Officer in conjunction with the Office of Human Resources Management (OHRM) has launched a department-wide AI learning initiative through its Office of Talent Programs. This initiative aims to provide every employee with both foundational and applied knowledge of AI, regardless of their prior technical experience by early 2026. The program focuses on cultivating AI literacy, reinforcing ethical standards, and encouraging practical exploration of AI use cases.

The training program is being deployed in two phases:

- Phase 1: This phase focuses on foundational AI literacy and responsible use. It includes AI Foundations Awareness assessment, two core courses (AI in the Workplace and Responsible Use of AI), and an optional AI Conversation Simulator tool.
- Phase 2: This phase offers advanced modules aligned with employees' roles and career paths, such as AI in Business Strategy, Machine Learning Fundamentals, AI and Data Privacy, and AI for Decision-Making.

This initiative supports the development of competencies critical to the Federal workforce's effective and ethical engagement with AI, including technology application, critical thinking, decision-making, AI and machine learning fundamentals, data literacy, and ethical AI implementation. The training components will be hosted on the Department's learning management system, the Commerce Learning Center (CLC), at no cost to employees. The curriculum will leverage CLC's content libraries and integrate GSA's AI Training Services to ensure alignment with cross-agency standards and best practices.

To sustain momentum and deepen AI capabilities, the Department will establish an internal AI Community of Practice (CoP) to support continuous learning, peer exchange, and innovation. Additional initiatives will include instructor-led workshops on advanced AI frameworks and integration of AI upskilling into performance plans and career development pathways.

Improving AI Governance

AI Governance Body

The Department of Commerce has established the Commerce AI Council (CAIC) as its primary agency-wide governance body for AI. The CAIC serves as the central advisory and oversight forum for AI strategy, operations, and compliance with Federal directives. Its mandate encompasses the review of agency AI priorities, the coordination of bureau-level efforts, and the provision of recommendations to the DOC CIO to ensure the responsible and effective use of AI across the Department.

Membership and Representation:
The CAIC is chaired by the DOC CIO and Chief AI Officer (CAIO) and features representation from all Bureau CIOs and their AI-specific support leads.

Bureau AI Leads are designated by Bureau CIOs and serve as the operational points of contact for AI-related activities, ensuring alignment with departmental standards and compliance obligations. The Council also draws upon support from technical subject-matter experts, program leads, and designated “Key Enabler” offices (e.g., enterprise architecture, data governance, acquisition, privacy, and security) as needed to support integrated governance.

Supporting Structures:
The CAIC is supported by the AI Integrated Project Team (AIIPT), a cross-functional body responsible for execution. The AIIPT delivers department-wide AI products and compliance deliverables, manages operational projects, and ensures implementation of CAIC guidance. Together, the CAIC and AIIPT provide a dual structure: the CAIC sets policy and strategic direction, while the AIIPT manages implementation and operationalization.

Mission and Objectives:
The CAIC is responsible for ensuring that AI adoption at DOC is consistent with Federal law, OMB memoranda, and Department policy. Its objectives include:

1. Strategic Alignment – Advise on AI priorities that support departmental and Federal missions.
2. Compliance & Ethics – Ensure that AI activities adhere to applicable statutes, regulations, and ethical guidelines, including privacy and civil rights protections.
3. Integration & Oversight – Oversee integration of AI capabilities into mission operations while ensuring security, accountability, and transparency.
4. Capacity Building – Promote workforce readiness through coordination of AI training, upskilling, and resource development.
5. Cross-Agency Coordination – Advance collaboration across bureaus and Federal agencies to maximize efficiency and alignment with government-wide AI initiatives.

Expected Outcomes:
Through this governance structure, DOC expects the following:

- Strengthen department-wide capacity to develop, test, and deploy AI solutions responsibly;
- Enhance transparency and accountability in AI adoption; and
- Achieve full compliance with OMB Memoranda M-25-21 and M-25-22;
- Standardize AI use case intake, review, and reporting across bureaus;
- Position the Department as a leader in advancing responsible and ethical Federal AI practices.

The Department of Commerce's AI governance process has full representation from across the agency including the following offices and bureaus:

Offices:

- Office of the Deputy Secretary
- Office of the Chief Information Officer
- Office of Cybersecurity and IT Risk Management Office of the General Counsel
- Office of the Chief Financial Officer and Assistant Secretary for Administration
- Office of Human Resources Management
- Office of Privacy & Open Government
- Office of Intelligence and Security
- Office of Civil Rights
- Office of Policy and Strategic Planning
- Office of the Chief Data Officer

Bureaus:

- Bureau of Economic Analysis
- Bureau of Industry and Security
- U.S. Census Bureau
- Economic Development Administration
- First Responder Network Authority
- International Trade Administration
- Minority Business Development Agency
- National Institute of Standards and Technology
- National Oceanic and Atmospheric Administration
- National Telecommunications and Information Administration
- National Technical Information Service
- Office of the Under Secretary for Economic Affairs

- U.S. Patent and Trademark Office

Agency Policies

The Department has undertaken a series of internal policy updates to align with the principles and directives outlined in OMB Memorandum M-25-21. These updates are designed to ensure that the Department's use of AI is consistent with federal mandates on innovation, risk management, and public trust.

Specifically, the Department has issued the AI Policy, approved by the DOC CIO. This policy establishes a governance framework that mandates:

- The development, acquisition, deployment, and use of AI technologies only when authorized and approved by the DOC's CIO. The documentation of all AI use cases in the centralized Commerce AI Use Case Inventory, with only DOC CIO-approved entries permitted for operational use.
- Continuous oversight of AI systems and environments, with the authority to revoke operational status (Authority to Operate, or ATO) for non-compliant systems. These measures are designed to ensure that AI deployments are secure, privacy-conscious, and aligned with broader IT infrastructure, cybersecurity, and data governance policies. The policy operates in concert with existing laws, regulations, and internal controls, including the Department's Rules of Behavior and cybersecurity protocols.

Internal Guidance on the use of Generative AI

DOC is actively developing internal guidance specific to the use of generative AI technologies in line with the directive and timing specified in OMB memorandum M-25-21, including the use of any OMB-provided template.

The current AI policy signed and in place at DOC is broader in scope than a specific GenAI policy as directed by OMB memorandum M-25-21 but its provisions apply broadly across all AI use cases, including those involving generative models.

In accordance with M-25-21, the Department is establishing safeguards and oversight mechanisms to mitigate risks associated with generative AI, including:

- Requiring CIO-level evaluation and approval of all generative AI tools prior to deployment.

- Mandating inclusion of generative AI use cases in the Commerce AI Use Case Inventory for transparency and accountability.
- Enforcing continuous monitoring and the potential revocation of ATO for any generative AI system that fails to comply with policy or poses undue risk to privacy, security, or mission integrity.

These efforts are informed by statutory authorities such as the AI in Government Act of 2020, the Advancing American AI Act, and Executive Orders 13960 and 14179, which collectively emphasize the importance of trustworthy, transparent, and secure AI in federal operations

AI Use Case Inventory

DOC has established an enterprise-wide intake and management process for soliciting, collecting, and maintaining artificial intelligence (AI) use cases across all bureaus, offices, and components. This process is coordinated by the OCIO in partnership with the AIIPT and CAIC, which collectively provide governance, oversight, and bureau-level accountability.

DOC's current process is to collect the required annual use case inventory within the timeframes prescribed by the Office of Management and Budget (OMB), establishing an annual baseline that is scrutinized against the previous year's submission, and accounting for the current year's update in guidance, reporting fields, and definitions.

To ensure comprehensive and consistent reporting, DOC has deployed a centralized AI Use Case Inventory platform, hosted on the Commerce AI Hub. This inventory is aligned to OMB-mandated reporting requirements and provides a single, authoritative system of record for AI activities. The platform enables 24/7 intake and maintenance of AI use cases, eliminating the reliance on periodic data calls and allowing bureaus to self-manage their portfolio of AI activities in real time.

Designated AI Integrated Project Team (AIIPT) Leads within each bureau are responsible for submitting, validating, and updating their bureau's entries. These Leads may delegate technical contributors as necessary but retain accountability for insuring that the information submitted is accurate, timely, and complete. The inventory encompasses all lifecycle stages of AI, including planned, pilot, deployed, and retired systems, and requires documentation on statutory data fields such as use case purpose, risk classification, vendor information, and applicable risk management controls.

DOC has instituted governance mechanisms to ensure that updates to existing use cases—including modifications to risk management documentation when high-impact AI enters into use—are systematically captured. Automated synchronization of bureau files with the enterprise inventory occurs twice daily through a Power Automate process, ensuring that any changes are reflected promptly in the master record. The AI Program team conducts routine quality assurance reviews and provides training and technical assistance to bureau Leads to ensure consistent compliance with reporting requirements.

The inventory is visible to all Commerce personnel to foster transparency and accountability, while write-access is restricted to authorized bureau representatives. Feedback mechanisms are embedded in the AI Hub interface, allowing bureaus to raise questions, identify gaps, and propose enhancements to the reporting process.

Through this structure, the Department ensures that its AI Use Case Inventory remains comprehensive, up to date, and reflective of all relevant activities across the Department, thereby satisfying statutory obligations, enhancing risk management, and supporting decision-making on the responsible use of AI.

Use case status is reported across the Department at the direction of the CAIO through its inclusion in the CAIC, and the DOC's CIO Council. Lastly, the CAIO coordinates use case reporting in other, non-IT governance groups such as the CDO Council, as necessary.

Fostering Public Trust in Federal Use of AI

Determinations of Presumed High-Impact AI

The Department of Commerce is aligning its internal AI governance framework with the definition of "high-impact AI" as articulated in Section 5 of the Appendix to OMB Memorandum M-25-21. While the Department has not yet issued a distinct internal definition, its current policy framework—codified in the Department's AI Policy, approved by the DOC CIO on September 19, 2025—establishes foundational controls that support the identification and oversight of high-impact AI systems.

Under the Department's AI policy:

- All AI use cases must be documented in the Commerce AI Use Case Inventory, which serves as the Department's centralized mechanism for tracking, reviewing, and approving AI deployments.
- Only AI systems that have been evaluated and approved by the DOC's CIO may be developed, acquired, or deployed.

- Continuous oversight is mandated for all AI systems, with the authority to revoke the Authority to Operate (ATO) for any system found to be non-compliant.

The Department is currently evaluating additional criteria to supplement the M-25-21 definition of high-impact AI.

These criteria will be formalized in forthcoming updates to the Department's AI governance documentation and integrated into the AI Use Case Inventory review process.

Waiving Minimum Risk Management Practices

The Department's CAIO, in appropriate coordination with other relevant officials, may waive one or more of the minimum risk requirements as outlined in OMB memorandum M-25-21 for a specific covered AI application or component. The waiver will be done so through a written determination, based upon a system-specific and context-specific risk assessment, that fulfilling the requirement would increase risks to safety or rights overall or would create an unacceptable impediment to critical agency operations. The Department CAIO will certify the ongoing validity of any waiver on an annual basis and reserves the right to revoke a previously issued waiver at any time. The Department CAIO will publicly report any determinations and waivers for AI use cases as described in OMB memorandum M-25-21 and any related follow-up instructions.

The Department CAIO plans to centrally track waivers, reassess if there are significant changes to the conditions or context in which the AI is used, and within 30 days of granting or revoking any waiver, report to 0MB on the scope, justification, and evidence supporting that action.

The Department is planning a formal waiver process to exempt AI use cases from the minimum risk management practices outlined in M-25-21 when required. The Department's AI Policy provides a governance structure that supports such a process, specifically:

- The DOC CIO retains sole authority to approve or deny AI use cases, which includes the discretion to impose or relax specific controls based on risk assessments.
- Any AI system not in compliance with policy requirements may have its ATO revoked, indicating a mechanism for enforcement and accountability.

As part of the Department's plans and assessment in developing a formal waiver issuance and tracking process, the following is considered as part of the process:

- A documented justification for the waiver request, including risk mitigation strategies.
- A review and approval workflow led by the DOC CIO or a designated AI governance board.
- A mechanism for tracking active waivers and their expiration or revocation status.

Implementation of Risk Management Practices and Termination of Non-Compliant AI

The Department AI policy explicitly authorizes the revocation of the Authority to Operate (ATO) for any AI system found to be in violation of departmental policy. Such systems may be shut down immediately to mitigate risk and ensure operational integrity.

Preventing Public Deployment of Non-Compliant High-Impact AI

All technology within DOC is subject to the U.S. Department of Commerce Technology Insertion (TI) Policy, which establishes how new technology products are reviewed and approved to promote compatibility, interoperability and conformance to statutory requirements as well as security and performance architecture prior to insertion into the DOC's IT ecosystem and inclusion into the Technology Standards List for use as IT assets within the DOC. As a part of its routine processes, the CAIO and supporting staff coordinates with the AI governance groups and the bureau CIOs to monitor the status of all reported use cases, pause the use of any potentially uncompliant AI, and terminate use cases, as necessary.

For AI specifically, the Department of Commerce has instituted a centralized governance framework to ensure that no high-impact AI system—defined in accordance with Section 5 of the Appendix to OMB Memorandum M-25-21—is deployed to the public without full compliance with applicable risk management practices.

Pursuant to the Department's Artificial Intelligence (AI) Policy, approved by the DOC CIO, the following controls and guidance have been put in place across the agency:

- Pre-Deployment Authorization: All AI systems, including those with potential high-impact characteristics, must undergo formal evaluation and receive explicit approval from the DOC's CIO prior to development, acquisition, or deployment.
- Centralized Inventory and Oversight: Each AI use case must be documented in the Commerce AI Use Case Inventory, which serves as the authoritative system of record for tracking and validating AI deployments. Only use cases marked as "DOC CIO approved" are authorized for use.
- Policy Enforcement and Legal Alignment: The AI Policy operates in conjunction with existing laws, regulations, and internal controls, including the Department's cybersecurity, privacy, and IT infrastructure policies. This ensures that AI systems are not only technically sound but also legally compliant.

These controls are designed to prevent the unauthorized release of AI systems that may pose undue risk to the public, particularly in cases where the AI system exercises autonomous decision-making or processes sensitive data.

Effective Termination of Non-Compliant AI

According to the Commerce IT policy, uncompliant systems may lose the authority to operate and may lead to the system being shut down until the violations have been remediated and the system is reauthorized. Potentially uncompliant AI is immediately paused, while the review and remediation timeframe prior to termination is set by the CAIO on a case-by-case basis. After the performance of an impact assessment, the CAIO's decision is reported through the Department's AI governance process.

The Department has established a clear and enforceable process for terminating AI systems that are found to be non-compliant with departmental policy or Federal risk management standards:

- Continuous Oversight: All AI systems are subject to ongoing monitoring and review. This includes periodic assessments to verify continued compliance with risk management, privacy, and security requirements1.
- Revocation of Authority to Operate (ATO): The Department's AI Policy grants the DOC CIO the authority to revoke the ATO for any AI system that is determined to be in violation of policy. Upon revocation, the system may be immediately shut down to mitigate operational, legal, or reputational risk1.
- Escalation and Remediation: In cases where non-compliance is identified, the Department may initiate an internal review to determine whether remediation is feasible. If not, the system will be decommissioned in accordance with established IT asset management and cybersecurity protocols.

This termination process ensures that the Department retains the ability to remove non-compliant AI systems from operation, thereby upholding public trust and minimizing institutional risk.

Issued By:

BRIAN EPLEY

Digitally signed by BRIAN EPLEY
Date: 2025.09.29 16:02:11 -04'00'

Brian Epley
Chief Information Officer (CIO) &
Chief AI Officer (CAIO)