U.S. Department of Commerce Compliance Plan for OMB Memoranda M-24-10 – September 2024.

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

1. STRENGTHENING AI GOVERNANCE

General

- Describe any planned or current efforts within your agency to update any existing internal AI principles, guidelines, or policy to ensure consistency with M-24-10.

  Artificial intelligence (AI) governance at the U.S. Department of Commerce (Commerce) centers on ensuring the coordination of key enablers for AI adoption and risk management across all its offices and bureaus, while promoting the centralization of AI approval and monitoring to the CAIO and its supporting units across the Office of the CIO. Ultimately, Commerce seeks to fully integrate AI governance within its enterprise information technology (IT) governance and management framework, treating AI as another value-add technology that is managed according to its statutory requirements, demand level, risk profile, and value proposition.

  The U.S. Department of Commerce Chief AI Officer is a participating member of the federal Chief AI Officer Council (CAIOC), and all Commerce AI governance is structured to follow and implement current and future CAIOC directives. Additionally, senior members of the Commerce Office of the CIO actively participate in all existing CAIOC working groups, with lead roles in a variety of tasks.

  Within Commerce, AI is subject to specialized governance, as well as broader IT oversight associated with cybersecurity, enterprise architecture, budget, customer experience, systems development lifecycle (SDLC), and other technology best practices. As mandated by EO13960, EO14110, and M-24-10, the Commerce CAIO holds final authority on AI governance, supported by key governance and support advisory bodies.

AI Governance Bodies

- Identify the offices that are represented on your agency's AI governance body.

  The Commerce AI Governance Board (CAIGB) – Serves as the highest, agency-wide governance body that coordinates and governs issues tied to the use of AI; chaired by the Deputy Secretary of Commerce, and vice chaired by the CAIO. The CAIGB reviews and comments on all policies, guidance, and directives issued by the CAIO, and serves as the final advisory resource for AI issues and decisions, featuring representation from senior agency officials that include IT, cybersecurity, data, privacy, civil rights and civil liberties, equity, statistics, human capital, procurement, budget, legal, agency management, customer experience, and program evaluation. The CAIGB is scheduled to meet

quarterly.

The Commerce AI Council (CAIC) – Serves as the primary agency-wide, advisory body for AI strategy, operations, and reporting at Commerce, and features representation from all bureau CIOs and their AI-specific support leads. The CAIC is chaired by the CAIO, and vice chaired by the Chief Technology Officer (CTO). As an operational council, the CAIC is supported by working groups and programs aligned with the M-24-10 requirements structure and focuses on practical outcomes. The CAIC is scheduled to meet monthly, or as needed.

The CAIC Governance and Operations Working Group (GOWG) – Operationalizes the requirements of M-24-10 Section 3 *Strengthening Artificial Intelligence Governance* and supports overall AI operations across Commerce.

The CAIC Advancing AI Working Group (AAWG) – Operationalizes the requirements of M-24-10 Section 4, *Advancing Responsible Artificial Intelligence Innovation*, and supports the promotion of mission-aligned and value-driven AI.

The CAIC Risk Management Working Group (RMWG) – Operationalizes the requirements of M-24-10 Section 5, *Managing Risks from the Use of Artificial Intelligence*, and supports the use of safe and ethical AI.

The Generative AI Pilot Review Group (PRG) – Operationalizes the governance of provisional use, generative AI systems for official purposes that advance the Department's mission and provides a rigorous evaluation of proposals based on seven framing principles that underlie the federal laws and guidance regarding AI. As articulated in the recent executive order and NIST's AI Risk Management, those principles are: (1) validity, reliability, and transparency; (2) safety and security; (3) innovation and competition; (4) worker support; (5) consideration of AI bias and civil rights; (6) consumer protection; and (7) privacy. The PRG is chaired by the AI Pilot Coordinator, with outcomes issued by the CAIO, with coordination with the Commerce Chief Operating Officer (COO).

Represented bureaus and offices within Commerce AI governance:

Offices:

- Office of the Deputy Secretary
- Office of the Chief Information Officer
- Office of Cyber Security and 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

- Describe the expected outcomes for the AI governance body and your agency's plan to achieve them.

  The expected governance outcomes vary by advisory body; however, all groups share the common goals of ensuring that all Commerce AI use has the approval and visibility of the CAIO and other necessary key enablers, that AI is promoted safely and wisely, that AI use supports Commerce's mission, and that Commerce remains compliant with industry and statutory guidance and best practices.

- Describe how, if at all, your agency's AI governance body plans to consult with external experts as appropriate and consistent with applicable law. External experts are characterized as individuals outside your agency, which may include individuals from other agencies, federally funded research and development centers, academic institutions, think tanks, industry, civil society, or labor unions.

  Commerce currently consults and coordinates with various external bodies and experts. Specifically, Commerce works with or regularly monitors:

  - The Chief AI Officer Council (CAIOC) working groups on
    - Procurement
    - Generative AI
    - Minimum Risk Practices
  - The National AI Advisory Committee (NAIAC)
  - The National Institute of Standards and Technology Information Technology Laboratory (ITL)
  - The National Institute of Standards and Technology AI Safety Institute
  - Other CFO Act agencies AI programs, as well as the Department of Defense

- Privately funded services that analyze or aggregate guidance, practices, or directives such as Gartner

AI Use Case Inventories

- Describe your agency's process for soliciting and collecting AI use cases across all sub-agencies, components, or bureaus for the inventory. *In particular, address how your agency plans to ensure your inventory is comprehensive, complete, and encompasses updates[4] to existing use cases.*

  The U.S. Department of Commerce currently owns one of the most comprehensive federal AI use case inventories publicly available. By including all bureau CIOs in the AI governance process, all use case owners are afforded the visibility of current and future requirements and are equipped to respond to changing requirements and updates. Commerce'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. Once the baseline inventory is published, bureau CIOs are responsible for the active maintenance of their use cases, and changes are flagged for review at the appropriate level of governance, based on the specifics of the use case. Inclusion of the use case into the Commerce AI inventory is a prerequisite to CAIO approval.

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

Reporting on AI Use Cases Not Subject to Inventory

- Describe your agency's process for soliciting and collecting AI use cases that meet the criteria for exclusion from being individually inventoried, as required by Section 3(a)(v) of M-24-10. *In particular, explain the process by which your agency determines whether a use case should be excluded from being individually inventoried and the criteria involved for such a determination.*

  Commerce follows the parameters within the *Guidance for 2024 Agency Artificial Intelligence Reporting per EO 14110*; each bureau is responsible for providing a unique rationale for exclusion at the time of the annual inventory request. There have not yet been any such requests, but the process of adjudication follows the standard AI governance escalation flow, where the CAIC and CAIGB provide insight for CAIO judgment. The primary consideration for non-reporting at this time is a privacy concern, where the simple publishing of a use case may expose an otherwise intentionally private process.

- Identify how your agency plans to periodically revisit and validate these use cases. *In particular, describe the criteria that your agency intends to use to determine whether an AI use case that previously met the exclusion criteria for individual inventorying should subsequently be added to the agency's public inventory.*

  When identified, approved excluded use cases will be flagged for annual review,

minimally. Additionally, the CAIO and CAIGB members have the ability to request status updates on use cases at any time.

2. ADVANCING RESPONSIBLE AI INNOVATION

Removing Barriers to the Responsible Use of AI

- Describe any barriers to the responsible use of AI that your agency has identified, as well as any steps your agency has taken (or plans to take) to mitigate or remove these identified barriers.[6] *In particular, elaborate on whether your agency is addressing access to the necessary software tools, open-source libraries, and deployment and monitoring capabilities to rapidly develop, test, and maintain AI applications.*

  AI governance remains a broadly unfunded requirement, severely impacting Commerce's ability to thoroughly analyze and track the responsible use of AI. To mitigate this issue, the CAIO is prioritizing funding requests for AI program support in the budget process, socializing the needs with the CAIGB, including senior leaders from across Commerce, and working to leverage the expertise and tools that already exists across the organization.

- Identify whether your agency has developed (or is in the process of developing) internal guidance for the use of generative AI. *In particular, elaborate on how your agency has established adequate safeguards and oversight mechanisms that allow generative AI to be used in the agency without posing undue risk.*

  On April 12, 2024, the Commerce CAIO issued departmental guidance on the approval, use, and monitoring of generative AI across the agency. Specifically, the guidance established Commerce's authority and purpose in regulating generative AI, described the criteria for evaluating and selecting generative AI proposals, outlined the overall process for proposing, approving, and monitoring generative AI pilot projects, and created the Generative AI Pilot Program. This guidance document is currently in effect across the agency but is also under consideration for integration into a greater Commerce AI policy.

AI Talent

- Describe any planned or in-progress initiatives from your agency to increase AI talent. *In particular, reference any hiring authorities that your agency is leveraging, describe any AI-focused teams that your agency is establishing or expanding, and identify the skillsets or skill-levels that your agency is looking to attract. If your agency has designated an AI Talent Lead, identify which office they are assigned to.*

  Commerce has not designated an AI Talent Lead, but efforts to increase AI talent are ongoing across the organization. At the operational unit level, bureaus typically evaluate AI expertise needs as part of the development of AI use cases and may include funds or resources in the acquisition phase. The Office of the CIO has also begun coordination with the Commerce Chief Learning Officer to increase organizational AI talent and expertise, and other complementary organizations, such as the Office of the CDO, have ongoing efforts to supplement AI talent within their specialties. Lastly, Commerce is also participating in government-wide efforts to increase AI talent, such as Service for America.

- If applicable, describe your agency's plans to provide any resources or training to develop AI talent internally and increase AI training opportunities for Federal employees. *In particular, reference any role-based AI training tracks that your agency is interested in, or actively working to develop (e.g., focusing on leadership, acquisition workforce, hiring teams, software engineers, administrative personnel or others).*

  Commerce considers effective AI training as an essential requirement for the promotion of mission-aligned and value-driven AI; staff are routinely allowed to attend AI trainings and conferences that augment their subject area's skills. AI training is readily available and encouraged across the organization; one recent example of agency-wide training available to all staff is the Commerce Research Library's Lexis CLE: *What is Legal Artificial Intelligence And Who is Regulating It*?

AI Sharing and Collaboration

- Describe your agency's process for ensuring that custom-developed AI code—including models and model weights—for AI applications in active use is shared consistent with Section 4(d) of M-24-10.

  Due to limited funding across the organization, Commerce presently focuses on the use of preferred, readily available solutions such as its Microsoft ecosystem to manage and share data. For specific AI applications, the Commerce Office of the CIO and the CDO are researching commercial-off-the-shelf solutions and repositories.

- Elaborate on your agency's efforts to encourage or incentivize the sharing of code, models, and data with the public. Include a description of the relevant offices that are responsible for coordinating this work.

  Commerce does not yet have an AI public dissemination roadmap; future implementation will be led by the CAIO and its CIO supporting offices, with support from the CDO, the Office of Privacy & Open Government, and the Office of Policy and Strategic Planning. Ultimately, all dissemination processes will be issued by the CAIO, with advice from the governing bodies.

Harmonization of Artificial Intelligence Requirements

- Explain any steps your agency has taken to document and share best practices regarding AI governance, innovation, or risk management. *Identify how these resources are shared and maintained across the agency.*

  Commerce routinely participates in external groups tasked with AI research and the issuance of guidance and best practices, and intents on continuing to import such practices. Also, some of the most significant experts in AI best practices, such as the National Institute of Standards and Technology Information Technology Laboratory (ITL) and the National Institute of Standards and Technology AI Safety Institute, reside within Commerce and are an integral part of the agency's AI governance process.

  Commerce AI governance is managed by the Office of the CTO, on behalf of the CAIO, and has access to a variety of cross-cutting communities and supplemental groups, such

as the Enterprise Architecture Program, allowing for multi-prong data retrieval and best practices dissemination.

3. MANAGING RISKS FROM THE USE OF ARTIFICIAL INTELLIGENCE

Determining Which Artificial Intelligence Is Presumed to Be Safety-Impacting or Rights-Impacting

- Explain the process by which your agency determines which AI use cases are rights-impacting or safety-impacting. *In particular, describe how your agency is reviewing or planning to review each current and planned use of AI to assess whether it matches the definition of safety-impacting AI or rights-impacting AI, as defined in Section 6 of M-24-10. Identify whether your agency has created additional criteria for when an AI use is safety-impacting or rights-impacting and describe such supplementary criteria.*

  Commerce follows the definitions of "Rights-Impacting AI" and "Safety-Impacting AI," as well as *Appendix I: Purposes for Which AI is Presumed to be Safety-Impacting and Rights Impacting* provided within M-24-10. The CAIO and its CIO supporting staff facilitate the dissemination of these parameters and provide oversight to the accuracy of the analysis, but the initial determination is a responsibility allocated to the use case owners during the reporting to the use case inventory. Procedurally, disagreements on the impact status are adjudicated through the standard AI governance escalation flow, where the CAIC and CAIGB provide insight for CAIO judgment.

- If your agency has developed its own distinct criteria to guide a decision to waive one or more of the minimum risk management practices for a particular use case, describe the criteria.

  Commerce has not developed any minimum risk management practices waive criteria.

- Describe your agency's process for issuing, denying, revoking, tracking, and certifying waivers for one or more of the minimum risk management practices.

  Commerce has not developed any minimum risk management practices waive criteria.

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

- Elaborate on the controls your agency has put in place to prevent non-compliant safety-impacting or rights-impacting AI from being deployed to the public.

  All technology within Commerce 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 Department of Commerce's (DOC) 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 its CIO 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.

- Describe your agency's intended process to terminate, and effectuate that termination of, any non-compliant AI.

  According to the Commerce TI 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 both the CAIGB and the CIO Council.

Minimum Risk Management Practices

- Identify how your agency plans to document and validate implementation of the minimum risk management practices. *In addition, discuss how your agency assigns responsibility for the implementation and oversight of these requirements.*

  While minimum risk practices execution is the responsibility of the use case owner, validation is a responsibility centralized within the CAIO and its CIO supporting staff, and is performed by leveraging the various working groups, and with advisement of the CAIGB and the CAIC.

Issued By:

BRIAN EPLEY

Digitally signed by BRIAN EPLEY
Date: 2024.10.04 13:43:53 -04'00'

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