THE DEPUTY SECRETARY OF TRANSPORTATION
WASHINGTON, DC 20590

September 24, 2024

MEMORANDUM FOR THE OFFICE OF MANAGEMENT AND BUDGET (OMB)

ISSUED BY: Polly Trottenberg
Deputy Secretary of Transportation

PREPARED BY: Mike Horton, DBA
Acting Chief Artificial Intelligence Officer

SUBJECT: US Department of Transportation (DOT) Compliance Plan for OMB Memorandum M-24-10 (September 2024)

This Compliance Plan conveys DOT's approach to achieving consistency with OMB Memorandum M-24-10 *Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence*. The plan aligns with M-24-10's three main pillars of Strengthening AI Governance, Advancing Responsible AI Innovation, and Managing Risks from AI.

The Department will execute this Compliance Plan commensurate with available resources and update the Plan as our understanding, experience, and Federal guidance mature. This Compliance Plan is applicable to all DOT Operating Administrations and Secretarial Offices only to the extent that it is consistent with the expressed language contained in 49 U.S.C. 106 and 40110 as applicable to the Federal Aviation Administration and Office of Inspector General.

1. STRENGTHENING AI GOVERNANCE

a. General. Internal AI stakeholders in the Office of the Secretary and the Operating Administrations will create and update Departmental AI-related principles, guidelines, and policies to align with this Compliance Plan to include the:

- AI Strategic Plan,
- AI Minimum Risk Management for Safety-Impacting and Rights-Impacting Use Cases,
- AI Governance Structure,
- IT Privacy, Records Management, Cybersecurity, and Data policy and guidance,
- ITIM 2023-005 AI Use Case Inventory Policy, and
- Generative AI Use Guidance.

b. AI Governance Body. DOT's Non-Traditional and Emerging Transportation Technology (NETT) Council serves as the Department's AI Governance Board. The Council comprises the Secretary (ex officio), Deputy Secretary (chair), Under Secretary of Transportation for Policy (vice chair), and other senior Departmental leaders.

The NETT Council was established in December 2018 as an internal DOT vetting body for new and emerging transportation technologies that are not yet established enough to fit into obvious modal categories or require new policy approaches.

In April 2024, DOT initiated updates to the Council's membership and charter to reflect its additional role of serving as the Department's AI Governance Board to govern the use of AI, remove barriers to the use of AI, and manage its inherent risks. The NETT Council first met in its capacity as DOT's AI Governance Board in April 2024 and will continue meeting in this capacity at least twice annually with those meeting chaired by the Deputy Secretary and vice-chaired by the Chief Artificial Intelligence Officer (CAIO). DOT will also continue convening the NETT Council outside of the AI Governance Board capacity to satisfy other needs and requirements, as appropriate.

The NETT Council will fulfill its role as DOT's AI Governance Board through the following functions:

- Review and approve all AI governance structures, processes, policies, and guidance.
- Approve criteria for the CAIO's exclusion of AI use cases from dissemination in the Public Use Case Inventory, including mission-sensitive, safety-sensitive, and exploratory, experimental, and unvalidated research scenarios.
- Approve criteria for designation of operational AI use cases as safety-impacting or rights-impacting.
- Incorporate external expert viewpoints to broaden the Council's perspective on technical, ethics, civil rights, workforce, and transportation-sector specific AI challenges, implications, and best practices in accordance with applicable law. These external collaboration efforts currently include the Department's Advanced Research Projects Agency – Infrastructure (ARPA-I) Request for Information (RFI) for Opportunities and Challenges of Artificial Intelligence (AI) in Transportation, released on May 3, 2024.
- Govern the establishment and providing oversight for the NETT Council AI Coordination and Activities Working Group and the Safety, Rights, and Security Review Advisory Committee.

NETT Council AI Coordination and Activities Working Group (AICA Working Group)

The NETT Council established the AI Coordination and Activities Working Group to support the CAIO in meeting M-24-10 coordination and AI activity tracking requirements and contribute to collaboration on AI compliance, governance, and guidance documents. This working group is chaired by the CAIO and vice-chaired by representatives from the Office of the Assistant Secretary for Transportation Policy (OST-P) and the Office of the Assistant Secretary for Research and Technology (OST-R). Members include representation from the Office of the Chief Information Officer (OCIO) responsible for cybersecurity, IT infrastructure, privacy, and records management; other areas within the Office of the Secretary of Transportation responsible for civil rights, human resources, budget, and training; the Office of General Counsel; Operating Administration (OA) representatives responsible for accelerating and helping guide AI development and safe adoption; and representatives from other research and statistical initiatives.

AI Safety, Rights, and Security Review Advisory Committee (SR2 Committee)

The NETT Council also established the SR2 Committee to assist the CAIO in reviewing and approving the operational deployment of all safety-impacting and rights-impacting AI use cases. The SR2 Committee is also responsible for performing the Security Review required by Executive Order 14110 Section 4.7(a) before AI data, custom code, and models are shared with the public. This Committee is chaired by the CAIO. Committee members include representatives from the Office of the Secretary (OST) responsible for safety, security, civil rights, privacy, and data, and relevant representatives from the sponsoring AI use case OA.

c. AI Acceleration Infrastructure. The CAIO, in collaboration with the OCIO and AICA Working Group, will accelerate safe, secure, and equitable AI development through leading-edge and best-practice compliance tools, technology, procedures, and education that will harness AI developer creativity, reduce barriers, and ensure continual risk management and compliance, especially for safety-impacting and rights-impacting use cases.

The *AI Accelerator Roadmap* provides OA developers with expedited access to a secure laboratory environment for experimentation with potential solutions to critical operational challenges and to robust development, testing, and deployment environments. The AI Accelerator Roadmap also provides researchers and their external research partners with their own secure access to approved state-of-the-art AI functionality.

The *AI Support and Collaboration Center* provides technical and non-technical employees with tools to educate themselves about AI capabilities and risks, get inspired through best-in-class Government and private sector use cases, collaborate with AI workgroups and communities of practice, partner with Department subject matter experts, and follow the AI Accelerator Roadmap to turn their AI concept into operational reality.

AI Accelerator Roadmap

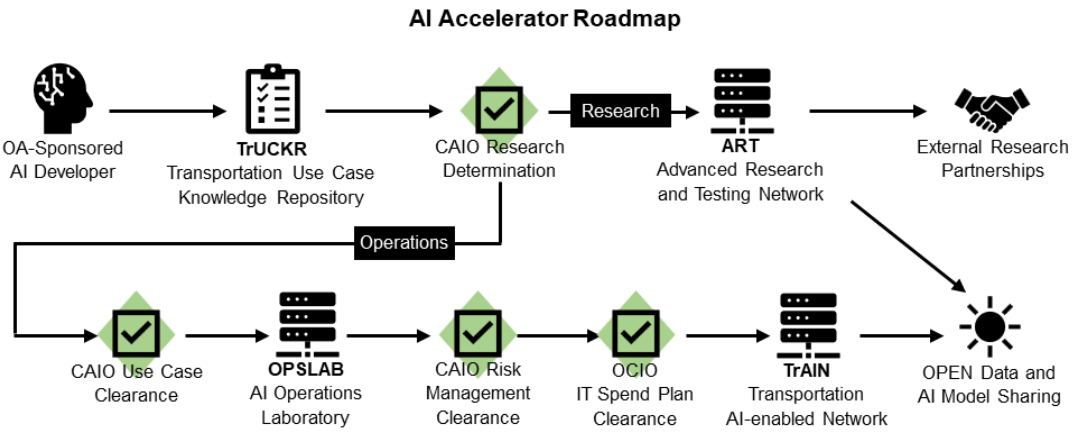

i. **Transportation Use Case Knowledge Repository (TrUCKR).** TrUCKR is the CAIO platform for tracking the Department's AI use case development, maturity, assessments,

clearances, risk evaluations and mitigations, and authorities to operate from conception through retirement.

Operational use case development, maturity, and risk compliance are driven and sponsored by each OA and Secretarial Office in coordination with the CAIO. Research use case development, maturity, and risk compliance are driven and sponsored by OST-R.

Operational use case AI Developers begin their AI Accelerator Roadmap journey by working with their OA to mature the use case for CAIO review, potential inclusion in the Public Use Case Inventory, and approval for initial concept development in the AI Operations Laboratory (OPSLAB).

Developers for use cases designated by the CAIO as "research" based on the requirements of OMB M-24-10 and other Federal Use Case Inventory guidance are directed to OST-R for further development in the Advanced Research and Testing Network as further discussed in the next clause. These research-based use cases are maintained in TrUCKR within the Restricted Use Case Inventory to fulfill OMB reporting requirements under M-24-10 for those use cases.

ii. **Advanced Research and Testing (ART) Network.** The ART Network is the OST-R-controlled and funded IT environments for AI research and development activities. It allows for rapid AI innovation, exploration, and sharing with external research partners while adhering to OCIO system requirements and CAIO compliance and risk mitigation mandates. A CAIO-led accelerated Authority to Operate (ATO) and system implementation process for emerging technology will provide researchers with essential leading-edge AI capabilities required by M-24-10 Section 4(b).

iii. **AI Operations Laboratory (OPSLAB).** OPSLAB is the CAIO-managed IT environments operated by the OCIO that is segmented from the rest of the Department's IT infrastructure. OPSLAB provides AI Developers with access to all OCIO-cleared AI functionality for use case experimentation, development and initial data and model risk management identification and mitigation. OPSLAB's primary purpose is to accelerate the determination of the required AI architecture, gain initial CAIO and SR2 Committee use case Authorization to Operate (ATO), and prepare for OCIO IT Spend Plan clearance and funding. All OPSLAB activity is managed and funded by the sponsoring OA.

iv. **Transportation AI-enabled Network (TrAIN).** The TrAIN supports the rapid deployment of AI solutions by aggregating all DOT AI-enabled development, testing, and production operational environments under the CAIO compliance and risk management monitoring umbrella where continuous AI model, AI data, and risk management monitoring will occur.

OA and Secretarial Offices can either opt to rapidly develop and deploy their AI solutions within a dedicated OCIO-managed environment or to create new, separate OA operational environments with the required CAIO compliance and risk management functionality by using established OCIO Authority to Operate processes.

v. **OPEN Data and AI Model Sharing.** The CAIO will ensure that operational use case AI data and models that receive SR2 Committee clearance will be shared with the public through established OPEN Data and Code.gov workflows.

vi. **Use Case Compliance Monitoring.** The CAIO, in collaboration with OAs and OST-R, will monitor all AI-enabled environments in the ART Network, OPSLAB, and TrAIN, as well as software that continually reports AI model purpose, impacts, and data usage. OAs are responsible for using these reports to enforce use case compliance and documented inclusion in TrUCKR. The CAIO will work to ensure that the OAs properly resolve use case application, component, and data usage discrepancies.

AI Support and Collaboration Center (AISCC)

The AISCC is the joint initiative of the CAIO, the Office of Innovation and Engagement (OIE), OST-R, and OCIO. Its mission is to accelerate safe, secure, transformative, and innovative AI solutions and research in DOT through employee education, collaboration, and governance through the following resources:

vii. **Get Educated.** *Learn.* Provides education videos, links, documents, information, and training on AI concepts and methodologies, as well as risk management considerations and mitigation approaches.

viii. **Get Inspired.** *Dream.* Provides access to the Department's Public Use Case Inventory, highlights use case lessons learned and best practices, and links to other Federal and private sector success stories.

ix. **Join a Community.** *Collaborate.* Provides listings and contact information for Department, Federal, and private sector AI workgroups and communities and their meeting artifacts where available.

x. **Find a Subject Matter Expert.** *Partner.* Provides the Department's subject matter expert listing, contact information, and areas of AI expertise.

xi. **Follow the AI Accelerator Roadmap Implementation Process.** *Make it Happen.* Provides detailed instructions, expectations, and examples for following the AI Accelerator Roadmap.

xii. **ASK AISCC.** *Get Answers.* Provides a tool for requesting support in conceptualizing and operationalizing use cases.

d. Public AI Use Case Inventories. The CAIO, in collaboration with OAs and OST-R, will ensure that all AI use cases will be tracked within TrUCKR through the AI Accelerator Roadmap from conception to retirement. TrUCKR will include comprehensive and complete use case status and compliance information, including the use case lifecycle stage, designation of rights-impacting or safety-impacting, documentation for risk management activities, and continuous authority to operate evaluations.

Through the AI Accelerator Roadmap, the CAIO ensures accountability for all AI use cases by utilizing TrUCKR as the entry point for access to all Departmental AI environments and continuous monitoring within the ART Network and TrAIN to identify AI models and associated data that are not authorized to operate.

The CAIO will also ensure that all AI use cases that meet the reporting requirements of Executive Order 13960 *Promoting the Use of Trustworthy Artificial Intelligence in Federal*

Government are identified within TrUCKR and reported in compliance with that order. The CAIO will also update the Department's *ITIM 2023-005: US Department of Transportation (DOT) Artificial Intelligence Use Cases Governance* to implement this process and align with the recent final OMB guidance on reporting requirements.

e. Reporting on AI Use Cases Not Subject to Public Inventory. The Department's AI Governance Board (NETT Council) sets the criteria for complying with Executive Order 13960 *Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government*. Based on those criteria, the CAIO effectuates determinations for exclusion from the Public Use Case Inventory in TrUCKR.

i. **Exclusion Process.** The CAIO applies the following exclusion criteria compliant with Executive Order 13960 to determine whether to include a use case in the Public Use Case Inventory before granting operational use case access to TrAIN:

   A. classified or sensitive,

   B. used in defense or national security systems as defined in 44 USC 2552(b)(6) (generally not applicable to DOT),

   C. embedded within standard commercial products, or

   D. research and development activities that meet the definition of basic research or applied research in M-24-10, where "basic research" is experimental or theoretical work undertaken primarily to acquire new knowledge of the underlying foundations of phenomena and observable facts without a specific application towards processes or products in mind, and "applied research" is an original investigation undertaken to acquire new knowledge to determine how a specific practical aim or objective may be met.

   The CAIO, through collaboration with the SR2 Committee, will also exclude from the Public Use Case Inventory AI use cases designated as mission-sensitive, safety-sensitive, confidential, or otherwise potential targets for malicious interference.

ii. **Exclusion Reevaluation Process.** The CAIO, in collaboration with OAs and the SR2 Committee, will update and revalidate all AI use cases within TrUCKR when that use case is modified but no less frequently than annually. During that update and revalidation process, the OA, CAIO, and SR2 Committee will evaluate if the use case exemption to reporting in the Public Use Case Inventory continues to meet the exclusion criteria.

2. ADVANCING RESPONSIBLE AI INNOVATION

a. Removing Barriers to the Responsible Use of AI. The Department's AI Accelerator Roadmap and AISCC provide the foundational tools, systems, best practices, playbooks, resources and procedures to responsibly, safely, and securely enable AI innovation and development.

i. **Barrier Identification and Mitigation.** Accelerating responsible AI adoption requires vigilance in identifying structural, procedural, educational, and training barriers and reducing compliance and administrative friction during development, all while

controlling risks to safety, privacy, accessibility, civil rights, and other rights, and adherence to applicable laws, regulations, and policies.

IT Infrastructure

The AI Accelerator Roadmap depicts the Department's IT infrastructure to ensure AI developer access to software tools, open-source libraries, secure cloud storage, and deployment and monitoring capabilities necessary to rapidly develop, test, and maintain AI applications.

A. **Advanced Research and Testing (ART) Network.** M-24-10 Section 4(a)(v) mandates that the Department provide sufficient AI tools and capacity to support research and development work. The ART Network accelerates AI Developer access to those tools by creating a platform of established, stand-alone, IT-compliant, AI-enabled environments. Through a collaboration between OST-R, OCIO, and the CAIO, the ART Network receives a prioritized review of critical and emerging technology platforms in Authorizations to Operate and other release or oversight processes to conform with guidelines discussed in M-24-10 Section 4(b)(iii).

B. **AI Operations Laboratory (OPSLAB).** The OPSLAB mirrors the stand-alone ART Network but focuses instead on operational AI developers for use case experimentation and maturation. Unlike the ART Network, OPSLAB also supports accelerated model compliance, security, and risk management evaluations, initial data quality, representativeness, and bias assessments, and the buildout of required system architecture and costing parameters necessary for IT Spend Plan approval, operational environment deployment, and Authorization to Operate.

C. **Transportation AI-enabled Network (TrAIN).** The TrAIN aggregates all Department AI-enabled development, test, and production (DTP) operational environments under one umbrella. The primary environments in TrAIN, typically one for each cloud provider of AI services, are part of the DTP service. These environments accelerate operational use case deployment, simplify safety, security, and risk monitoring, and reduce administrative friction, costs, and delays created by the need to establish new environments for each use case along with the required AI model and data compliance and risk management surveillance tools.

Data

The AI Accelerator Roadmap provides the Department with adequate infrastructure to share, curate, and govern data used in training, testing, and operating AI. This infrastructure includes resources that enable sound data governance and management best practices, including assessing all AI-related datasets for quality, accuracy, functionality, representativeness, and bias.

All AI training datasets in the research-based ART Network, operations-based OPSLAB, and TrAIN will use commercially available tools and best practices to build trustworthiness and to continually monitor those data for unacceptable bias to help ensure fair, equitable, and inclusive results that appropriately balance model fairness, performance, and real-world impact.

Cybersecurity

The AI Accelerator Roadmap also provides researchers and OAs access to AI-enabled environments with established continuously monitored AI model Authorizations to Operate. Cybersecurity in these platforms is further enhanced by segregating ART Network and OPSLAB environments from the Department's operational IT infrastructure.

The Department is also developing the ability to prioritize Authorizations to Operate and other release authorizations for generative AI and other critical and emerging technologies to accelerate AI research, operational development, and adoption.

ii. **Generative AI Internal Guidance.** The CAIO, in collaboration with OST-P, OST-R, OCIO, and OAs, is establishing policy, safeguards, and oversight mechanisms that will embrace the benefits of GenAI while mitigating its risks.

b. AI Talent. The precursor to accelerated, safe, and secure AI adoption is the acquisition and maintenance of a well-educated and trained workforce empowered with the skills, resources, guidance, inspiration, creative freedom, and implementation roadmap to stay abreast of evolving AI capabilities and risks and contribute AI-enabled solutions that assist DOT in delivering the world's leading transportation system.

i. **AI Talent Acquisition.** DOT is looking to increase AI talent throughout the OAs, which includes using all applicable hiring authorities and flexibilities. DOT's Chief Artificial Intelligence Officer position is established and will serve as a central source of information to support other AI recruitment activities throughout the Department. In addition, DOT is providing support to the Office of Personnel Management in its efforts to better define AI for the Federal workforce including appropriate occupational series and duties to be used for AI positions and position titles. To promote greater understanding of AI and how it impacts workforce planning and hiring initiatives, a learning session was held for the DOT Human Resources community. This session provided an overview of AI, including AI concepts and terminology, and promoted the use of tagging vacancies on USA Jobs that are AI related to allow potential applicants to find the AI related vacancy announcements more easily.

ii. **Internal AI Training.** The AISCC is the centralized, self-service hub for promoting the development of AI talent internally, providing pathways to AI occupations, and assisting employees affected by the application of AI to their work. An executive learning session was held to level set executives understanding of AI and terminology based on EO 14110. In addition, AI Day showcased AI efforts within DOT and use cases throughout the OAs. The event was open to all DOT employees.

c. AI Sharing and Collaboration. DOT is committed to the open sharing of AI custom code, models, and data that promote the reuse and collaboration with the Federal Government and public to enhance innovation and transparency while maintaining the public's rights, safety, and security.

i. **Custom-Developed AI Code.** As the world leader in the transportation sector, the CAIO will prioritize the sharing of custom-developed code, including commonly used packages and functions, models, and model weights, which have potential for reuse by other agencies and the public to the maximum extent possible in compliance with M-24-10 Section 4(d)(i).

ii. **Public Sharing.** The CAIO will ensure TrAIN-related data, custom code, and models that clear the SR2 Committee's security review as required under Executive Order 14110 Section 4.7(a) are shared with the public. The CAIO will maintain use case security review justification, documentation, and sharing methodologies in TrUCKR.

The CAIO, in collaboration with the use case owners, will ensure cleared data are shared through established OPEN Data workflows. Custom code will be released through Code.gov using the guidance and best practices found in OMB Memorandum M-16-21, *Federal Source Code Policy: Achieving Efficiency, Transparency, and Innovation through Reusable and Open Source Software (August 8, 2016)*, collaboration methods found in Executive Order 14110, the General Services Administration's AI Community of Practice, and other Federal requirements to include the OST-R-led Public Access Plan for research data.

d. Harmonization of Artificial Intelligence Requirements. The CAIO along with other members of the DOT AI community are active collaborators with the Office of Science and Technology Policy (OSTP), the National Institute of Standards and Technology (NIST), the Chief Artificial Intelligence Officer Council (CAIOC), and other Federal entities that seek to interpret and implement AI management requirements consistently across Federal agencies and create efficiencies and opportunities for sharing resources and best practices.

3. MANAGING RISKS FROM THE USE OF ARTIFICIAL INTELLIGENCE

A clear risk management process and documentation system are essential for effectively governing AI risks and ensuring that all use cases are properly monitored. The AI Accelerator Roadmap ensures that all use cases are correctly assessed as early as possible in the AI lifecycle. The CAIO makes initial safety-impacting or rights-impacting risk use case determinations during the Use Case Clearance stage before authorization to access the AI Operations Laboratory. The CAIO re-evaluates that determination in the Risk Management Clearance stage before permission is granted for access to TrAIN for development and testing. CAIO and SR2 Committee clearance is required prior to use case advancement into production. Reassessment determinations are conducted at least annually or when significant use case changes are made.

The TrUCKR platform maintains CAIO context-specific and system-specific risk evaluation, reevaluation, determination, reassessment, certification, and reporting documentation for each use case. All initial decisions and any changes in determination prompted by significant

modifications to the conditions or context in which the AI is used, including the scope, justification, and supporting evidence, will be reported by the CAIO to OMB within 30 days of the decision. In conjunction with M-24-10 Section 5(a)(ii), the CAIO will ensure a summary of each use case determination and waiver, including its justification, will be publicly released as required by Executive Order 13960 Use Case Inventory guidance.

a. Determining Which Artificial Intelligence Is Presumed to Be Safety-Impacting or Rights-Impacting. When an underlying AI use case status may be safety-impacting or rights-impacting, the CAIO will assess whether the AI application or component output would serve as a principal basis for a decision or action that will be used in real-world conditions, or significantly influence the outcomes of Department activities or decisions that impact safety or rights. The CAIO will base these assessments on the advice of the SR2 Committee using criteria approved by DOT's AI Governance Board.

i. **Safety-Impacting and Rights-Impacting Determinations.** OAs are responsible for adequately identifying, evaluating, and continually monitoring each AI use case for its potential and realized impact on safety and rights and sufficiently documenting those assessments and reassessments in TrUCKR for CAIO initial determination. The NETT Council's SR2 Committee advises the CAIO on the final determination for each use case prior to deployment. OAs can appeal that the determination to the NETT Council.

   The CAIO uses *M-24-10 Appendix I: Purposes for Which AI is Presumed to be Safety-Impacting and Rights-Impacting* as the general criteria for making these determinations. Through the CAIO's annual review of the evolution of deployment context, risks, benefits, and needs, the CAIO will recommend additional criteria and requirements for approval by the NETT Council.

ii. **Agency-Developed Minimum Risk Management Practice Waiver Criteria.** The Department does not anticipate any agency-developed minimum risk criteria that would waive use case compliance with minimum risk management practices defined in M-24-10.

iii. **Waiver Processes.** OAs can request a waiver to minimum risk management requirements and provide supporting justification within TrUCKR for consideration by the CAIO and SR2 Committee. All CAIO final waiver decisions will be reviewed by the NETT Council and documented in TrUCKR.

b. Implementation of Risk Management Practices and Termination of Non-Compliant AI. OAs are responsible for following the AI Accelerator Roadmap, ensuring continuous use case compliance with any minimum risk management requirements for their safety-impacting and rights-impacting use cases throughout the AI lifecycle, and reporting in TrUCKR any changes in AI application or component impacts on safety or rights for CAIO reassessment.

i. **Non-Compliant Controls.** All research and operational development of AI applications and components will be implemented through the Department's AI Accelerator Roadmap and included in TrUCKR. All environments within that Roadmap, including the ART Network, OPSLAB, and TrAIN, will deploy AI model and data usage monitoring software to continuously evaluate use case and data usage compliance as well as

minimum risk management compliance for safety-impacting and rights-impacting use cases.

OAs are responsible for using these reports to enforce compliance. The CAIO is responsible for ensuring that the OA use case and risk monitoring and reporting responsibilities are met. No safety-impacting or rights-impacting use case will be authorized for deployment if it does not meet and maintain minimum risk management compliance.

ii. **Non-Compliant Termination.** All safety-impacting and rights-impacting AI applications and components in production will be under a continuous Authority to Operate (ATO) initially issued by the CAIO as part of the CAIO Risk Management Clearance action in the AI Accelerator Roadmap. To receive and maintain continuous AI-model ATO, safety-impacting and rights-impacting use cases will be required to maintain an alternative process that does not depend on the AI-enabled capability. Use cases out of compliance with minimum risk standards, as determined by the CAIO through advisement with the SR2 Committee, will suspend operations and revert to the non-AI process until compliance is reinstituted and the use case is cleared to resume operations by the CAIO or terminated if minimum risk standards cannot be met.

c. Minimum Risk Practices Through the AI Accelerator Roadmap and its associated processes and AI-enabled environments, DOT has the foundation for meeting, documenting, and governing the M-24-10 Section 5(c) minimum risk management practices as follows:

- TrUCKR will document use case adherence to minimum risk management tracking requirements throughout the AI use case lifecycle.
- OAs will document potential use case exposure to safety-impacting and rights-impacting risks within the initial use case entry into TrUCKR.
- During use case maturity in OPSLAB, use cases determined by the CAIO through consultation with the SR2 Committee as safety-impacting or rights-impacting will require the OAs to update TrUCKR with a detailed Minimum Risk Management Mitigation Plan to receive use case CAIO ATOs and begin AI use case development in TrAIN.
- Before AI application deployment, the OA must document the completion of the initial Minimum Risk Management Mitigation Plan in TrUCKR to receive CAIO and SR2 Committee use case ATO in TrAIN production.
- OAs will also need to certify and document continued use case adherence to the Minimum Risk Management Mitigation Plan annually or when significant changes to the conditions or context in which the AI is used to maintain AI use case ATO in TrAIN.

4. DEFINITIONS

This document incorporates the following relevant definitions as provided in OMB Memoranda M-24-10 *Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence*:

Artificial Intelligence (AI). The term "artificial intelligence" has the meaning provided in Section 238(g) of the John S. McCain National Defense Authorization Act for Fiscal Year 2019, which states that "the term 'artificial intelligence' includes the following":

1. Any artificial system that performs tasks under varying and unpredictable circumstances without significant human oversight, or that can learn from experience and improve performance when exposed to data sets.
2. An artificial system developed in computer software, physical hardware, or other context that solves tasks requiring human-like perception, cognition, planning, learning, communication, or physical action.
3. An artificial system designed to think or act like a human, including cognitive architectures and neural networks.
4. A set of techniques, including machine learning, that is designed to approximate a cognitive task.
5. An artificial system designed to act rationally, including an intelligent software agent or embodied robot that achieves goals using perception, planning, reasoning, learning, communicating, decision making, and acting.

For the purposes of this memorandum, the following technical context should guide interpretation of the definition above:

1. This definition of AI encompasses, but is not limited to, the AI technical subfields of machine learning (including deep learning as well as supervised, unsupervised, and semi-supervised approaches), reinforcement learning, transfer learning, and generative AI.
2. This definition of AI does not include robotic process automation or other systems whose behavior is defined only by human-defined rules or that learn solely by repeating an observed practice exactly as it was conducted.
3. For this definition, no system should be considered too simple to qualify as covered AI due to a lack of technical complexity (e.g., the smaller number of parameters in a model, the type of model, or the amount of data used for training purposes).
4. This definition includes systems that are fully autonomous, partially autonomous, and not autonomous, and it includes systems that operate both with and without human oversight.

AI Maturity. The term "AI maturity" refers to a Federal Government organization's capacity to successfully and responsibly adopt AI into their operations and decision-making across the organization, manage its risks, and comply with relevant Federal law, regulation, and policy on AI.

AI Model. The term "AI model" has the meaning provided in Section 3(c) of Executive Order 14110 that states "a component of an informational system that implements AI technology and uses computational, statistical, or machine-learning techniques to produce outputs from a given set of inputs".

Applied Research. The term "applied research" refers to original investigation undertaken in order to acquire new knowledge to determine the means by which a specific practical aim or objective may be met.

Automation Bias. The term "automation bias" refers to the propensity for humans to inordinately favor suggestions from automated decision-making systems and to ignore or fail to seek out contradictory information made without automation.

Basic Research. The term "basic research" refers to experimental or theoretical work undertaken primarily to acquire new knowledge of the underlying foundations of phenomena and observable facts without a specific application towards processes or products in mind.

Custom-Developed Code. The term "custom-developed code" has the meaning provided in Appendix A of OMB Memorandum M-16-21 which states "custom-developed code is code that is first produced in the performance of a Federal contract or is otherwise fully funded by the Federal Government. It includes code, or segregable portions of code, for which the Government could obtain unlimited rights under Federal Acquisition Regulations (FAR) Pt. 27 and relevant agency FAR Supplements. Custom-developed code also includes code developed by agency employees as part of their official duties. For the purposes of this policy, custom-developed code may include, but is not limited to, code written for software projects, modules, plugins, scripts, middleware, and APIs; it does not, however, include code that is truly exploratory or disposable in nature, such as that written by a developer experimenting with a new language or library.

Data Asset. The term "data asset" has the meaning provided in 44 USC § 3502 which states "the term 'data asset' means a collection of data elements or data sets that may be grouped together."

Equity. The term "equity" has the meaning provided in Section 10(a) of Executive Order 14091 that states "the term 'equity' means the consistent and systematic treatment of all individuals in a fair, just, and impartial manner, including individuals who belong to communities that often have been denied such treatment, such as Black, Latino, Indigenous and Native American, Asian American, Native Hawaiian, and Pacific Islander persons and other persons of color; members of religious minorities; women and girls; LGBTQI+ persons; persons with disabilities; persons who live in rural areas; persons who live in United States Territories; persons otherwise adversely affected by persistent poverty or inequality; and individuals who belong to multiple such communities."

Generative AI. The term "generative AI" has the meaning provided in Section 3(p) of Executive Order 14110 which states "the term 'generative AI' means the class of AI models that emulate the structure and characteristics of input data in order to generate derived synthetic content. This can include images, videos, audio, text, and other digital content."

Model Weight. The term "model weight" has the meaning provided in Section 3(u) of Executive Order 14110 that states "the term 'model weight' means a numerical parameter within an AI model that helps determine the model's outputs in response to inputs."

Open Government Data Asset. The term "open government data asset" has the meaning provided in 44 USC § 3502 that states "the term 'open Government data asset' means a public data asset that is machine-readable; available (or could be made available) in an open format; not encumbered by restrictions, other than intellectual property rights, including under titles 17 and 35, that would impede the use or reuse of such asset; and based on an underlying open standard that is maintained by a standards organization."

Open Source Software. The term "open source software" has the meaning provided in Appendix A of OMB Memorandum M-16-21 that states "Open Source Software (OSS) is software that can be accessed, used, modified, and shared by anyone. OSS is often distributed under licenses that

comply with the definition of "Open Source" provided by the Open Source Initiative (https://opensource.org/osd) and/or that meet the definition of "Free Software" provided by the Free Software Foundation (https://www.gnu.org/philosophy/free-sw.html).

Rights-Impacting AI. The term "rights-impacting AI" refers to AI whose output serves as a principal basis for a decision or action concerning a specific individual or entity that has a legal, material, binding, or similarly significant effect on that individual's or entity's:

1. Civil rights, civil liberties, or privacy, including but not limited to freedom of speech, voting, human autonomy, and protections from discrimination, excessive punishment, and unlawful surveillance;
2. Equal opportunities, including equitable access to education, housing, insurance, credit, employment, and other programs where civil rights and equal opportunity protections apply; or
3. Access to or the ability to apply for critical government resources or services, including healthcare, financial services, public housing, social services, transportation, and essential goods and services.

Risks from the Use of AI. The term "risks from the use of AI" refers to risks related to efficacy, safety, equity, fairness, transparency, accountability, appropriateness, or lawfulness of a decision or action resulting from the use of AI to inform, influence, decide, or execute that decision or action. This includes such risks regardless of whether:

1. The AI merely informs the decision or action, partially automates it, or fully automates it;
2. There is or is not human oversight for the decision or action;
3. It is or is not easily apparent that a decision or action took place, such as when an AI application performs a background task or silently declines to take an action; or
4. The humans involved in making the decision or action or that are affected by it are or are not aware of how or to what extent the AI influenced or automated the decision or action.

While the particular forms of these risks continue to evolve, at least the following factors can create, contribute to, or exacerbate these risks:

1. AI outputs that are inaccurate or misleading;
2. AI outputs that are unreliable, ineffective, or not robust;
3. AI outputs that are discriminatory or have a discriminatory effect;
4. AI outputs that contribute to actions or decisions resulting in harmful or unsafe outcomes, including AI outputs that lower the barrier for people to take intentional and harmful actions;
5. AI being used for tasks to which it is poorly suited or being inappropriately repurposed in a context for which it was not intended;
6. AI being used in a context in which affected people have a reasonable expectation that a human is or should be primarily responsible for a decision or action; and
7. the adversarial evasion or manipulation of AI, such as an entity purposefully inducing AI to misclassify an input.

This definition applies to risks specifically arising from using AI and that affect the outcomes of decisions or actions. It does not include all risks associated with AI, such as risks related to the privacy, security, and confidentiality of the data used to train AI or used as inputs to AI models.

Safety-Impacting AI. The term "safety-impacting AI" refers to AI whose output produces an action or serves as a principal basis for a decision that has the potential to significantly impact the safety of:

1. Human life or well-being, including loss of life, serious injury, bodily harm, biological or chemical harms, occupational hazards, harassment or abuse, or mental health, including both individual and community aspects of these harms;
2. Climate or environment, including irreversible or significant environmental damage;
3. Critical infrastructure, including the critical infrastructure sectors defined in Presidential Policy Directive 21 or any successor directive and the infrastructure for voting and protecting the integrity of elections; or,
4. Strategic assets or resources, including high-value property and information marked as sensitive or classified by the Federal Government.

Significant Modification. The term "significant modification" refers to an update to an AI application or to the conditions or context in which it is used that meaningfully alters the AI's impact on rights or safety, such as through changing its functionality, underlying structure, or performance such that prior evaluations, training, or documentation become misleading to users, overseers, or individuals affected by the system. This includes significantly changing the context, scope, or intended purpose in which the AI is used.