

PART III: DOT COMPLIANCE PLAN FOR OMB MEMORANDUM M-25-21 (SEPTEMBER 2025)

This document presents the U.S. Department of Transportation's (DOTs) formal Compliance Plan in response to the Office of Management and Budget (OMB) Memorandum M-25-21, *Accelerating Federal Use of AI through Innovation, Governance, and Public Trust*. This is a living document and will be updated as DOT continues to evolve through the IDOT technology modernization initiative. The Department fully embraces the memorandum's pro-innovation, America-first approach to Artificial Intelligence (AI) as a critical enabler of its mission to advance safety, rebuild our national infrastructure, and usher in a new era of American transportation leadership. This plan serves as a direct operational component of the broader DOT AI Strategy, translating the strategic vision for leveraging AI in Safety, Infrastructure, Innovation, and Efficiency into a concrete, actionable framework for execution. It provides the "how" for the ambitious "what" and "why" articulated in the Department's strategic vision, ensuring that a portfolio of transformative AI Use Cases become operational realities.

The Department's core philosophy is one of responsible acceleration to aggressively remove bureaucratic barriers to empower our Nation's transportation innovators while maintaining the streamlined, effective governance necessary to ensure AI is deployed safely, fosters public trust, and maximizes taxpayer value. This includes ensuring appropriate safeguards are in place to protect privacy, civil rights, and civil liberties, and to mitigate any unlawful discrimination. The approach directly aligns with M-25-21's three key priorities of innovation, governance, and public trust. The structures detailed within this Plan are not designed to be bureaucratic constraints but are strategic enablers of accelerated, responsible innovation. By providing clear, predictable "guardrails," US DOT gives developers and program managers the confidence to move with speed and purpose. This proactive posture is informed by the Department's comprehensive AI Maturity Assessment, which provides an evidence-based analysis of our current capabilities. The assessment reveals a common challenge in large-scale technology adoption: while DOT has established a strong strategic and governance foundation, a gap exists in the core technical and enabling capabilities required for execution at scale. This Compliance Plan helps to close that execution gap, ensuring our visionary use cases become operational realities.

The Department will execute this Compliance Plan commensurate with available resources and will update the Compliance Plan every two years as our understanding, experience, and Federal guidance mature.

DRIVING INNOVATION

The U.S. Department of Transportation is committed to a proactive, enterprise-wide strategy for driving AI innovation, directly responding to OMB Memorandum M-25-21's mandate to "remove barriers to innovation" and "lean forward on adopting effective, mission-enabling AI to benefit the American people" while ensuring appropriate safeguards are in place to protect privacy, civil rights, and civil liberties, and to mitigate any unlawful discrimination. The Department is moving decisively beyond ad-hoc experimentation toward a deliberate, centrally managed ecosystem designed to foster and accelerate the responsible adoption of AI. Our approach recognizes that true innovation is not born from technology alone, but from a holistic system that provides innovators with the tools, data, and talent they need to succeed. This system is built on three core pillars that address the entire innovation lifecycle. The first pillar is Providing the Tools for Acceleration. The Department has already created its AI Accelerator Roadmap that will implement technologies to provide developers with secure, pre-configured environments for research, development, and deployment. This Roadmap is a strategic asset that slashes the administrative and security-related friction that traditionally grinds government innovation to a halt, enabling our teams to move from concept to operational reality at the speed of relevance.

The second pillar is Maximizing Taxpayer Value Through Reuse. To combat the inefficiency of siloed development, the Department is establishing a robust framework for sharing and reusing AI code, models, and data assets across its Operating Administrations (OAs). Led by the acting Chief Data and Artificial Intelligence Officer (CDAIO), this initiative proposes to establish a "share by default" culture supported by modern technical infrastructure, including a central source code repository and a library of curated, AI-ready datasets. By preventing the constant reinvention of foundational capabilities, DOT ensures that every taxpayer dollar invested in AI is maximized, allowing new projects to build upon prior successes rather than starting from scratch. This approach not only accelerates development timelines but also improves the quality and consistency of AI solutions across the Department.

The third and most critical pillar is Building a World-Class Workforce. Technology and infrastructure are inert without the skilled human capital to leverage them. The Department plans to execute a comprehensive plan to recruit, train, and empower the AI talent necessary to maintain America's global dominance in transportation technology. This effort involves an aggressive strategy to attract top technical talent from the private sector by utilizing all available hiring flexibilities, as well as a deep commitment to "upskilling" our existing workforce. The AI Support and Collaboration Center (AISCC) will serve as the hub for this effort, providing a wealth of self-service training resources, expert support, and communities of practice to foster a culture of continuous learning and responsible innovation. Together, these three pillars demonstrate that DOT has not only a vision for innovation

but will also have the concrete infrastructure, governance, and human capital plans required to make that vision a reality, ensuring full alignment with the spirit and letter of M-25-21.

REMOVING BARRIERS TO THE RESPONSIBLE USE OF AI

The Department's AI Maturity Assessment identified the primary barriers to responsible AI use not as a lack of innovative ideas, but as structural impediments to execution. These include fragmented access to computing resources, inconsistent development tools, lack of curated, AI-ready data, and security and compliance processes that often create friction and lead to duplicative efforts across OAs. To mitigate and remove these barriers, DOT's primary strategy is the AI Accelerator Roadmap, a comprehensive plan for an ecosystem designed to provide a streamlined, secure, and efficient "path to production" for AI developers. This roadmap directly addresses the M-25-21 requirement to provide AI-enabling infrastructure across the AI lifecycle including development, testing, deployment, and continuous monitoring.

The AI Accelerator Roadmap addresses access to necessary software tools, open-source libraries, and development capabilities through three specialized environments:

- **AI Operations Laboratory (OPSLAB):** This laboratory is designed for the experimentation and maturation of operational use cases, removing the lengthy procurement and security approval processes that typically stall projects at their inception.
- **Advanced Research and Testing (ART) Network:** This separate environment managed by the Office of the Assistant Secretary for Research and Technology provides researchers and their external partners with access to leading-edge, state-of-the-art AI functionality. ART directly fulfills the M-25-21 mandate to "provide AI tools and capacity to support the agency's AI research and development (R&D) efforts".
- **Transportation AI-enabled Network (TrAIN):** TrAIN provides pre-configured, compliant environments with built-in continuous monitoring capabilities. TrAIN dramatically accelerates the deployment timeline while ensuring robust oversight of all AI models and data in operation. This structure operationalizes the implementation plan for Pillar 4 (AI & Data Tech Infrastructure) from the AI Maturity Assessment, transforming a strategic goal into a tangible capability.

SHARING AND REUSE

To promote the sharing and reuse of AI code, models, and data assets, DOT's internal coordination is formally led by the CDAIO, as encouraged by M-25-21, through the Safety, Rights, and Security Review (SR2) Committee focused on accelerating and governing appropriate asset sharing. This group includes representatives from the

Office of the Chief Information Officer (OCIO) and Privacy, Recordkeeping, Safety, Civil Rights and Security leaders.

The Department is promoting a culture of sharing through both policy and technology:

- **Policy Mandate:** As part of its maturity assessment, DOT anticipates updating its existing “share by default” internal data and software policies to ensure they apply to AI-ready data and models and its “design for reuse” principle for all new AI development. This practice aligns directly with the implementation plan for Pillar 9 (AI Data, Model, & Code Sharing) of the AI Maturity Assessment.
- **Technical Infrastructure:** The Department plans to establish a central internal source code repository for all AI projects and a library of reusable models and code components. For data assets, the CDAIO is assisting with the creation of curated, AI-ready datasets for common use cases, which will be made discoverable through an enterprise data catalog, operationalizing key goals from the AI Maturity Assessment.

AI TALENT

DOT anticipates executing a two-pronged strategy to enhance AI talent: aggressively “upskilling” the existing workforce and strategically recruiting new talent. The primary vehicle for internal workforce development is the AI Support and Collaboration Center (AISCC). The “Get Educated” and “Join a Community” functions of the AISCC will provide a centralized, self-service hub for training videos, best practice guides, and connections to communities of practice, directly addressing the M-25-21 call to “leverage AI trainings and resources to upskill existing staff.” For recruitment, the CDAIO is partnering with the Office of Human Resources to leverage all available hiring authorities and flexibilities, potentially including the use of skills-based assessments and participation in government-wide AI talent surges, as recommended in M-25-21.

The Department's AI Maturity Assessment identifies critical skillsets needed at DOT, including AI/ML Engineers, Data Scientists, and Cloud Financial Operations (FinOps) specialists. Individuals with this technical talent can have the most impact in central offices like the OCIO and the Office of the CDAIO, where they will build and manage the enterprise platforms (TrAIN, OPSLAB, ART Network). Those with these skills are also critically needed within high-priority OA programs.

To develop the DOT workforce in these AI and AI-enabling roles, the Department plans to provide several key resources and training initiatives. In alignment with the Maturity Assessment, DOT intends to formalize an internal skills certification program to recognize acquired AI skills, creating clear pathways for internal career progression and mobility. The Department is also evaluating its ability to launch a formal mentorship program that pairs senior AI experts with emerging talent to cultivate a culture of responsibility and technical excellence. Finally, the AISCC's “ASK AISCC” function hopes to expand from a simple help desk to a proactive support

service, providing direct, hands-on assistance for employees in the future who are seeking to apply AI to their work, thereby moving from passive education to active project enablement.

IMPROVING AI GOVERNANCE

The U.S. Department of Transportation has established a governance framework that is not a bureaucratic constraint but a strategic enabler of accelerated, responsible innovation. This approach is fully consistent with OMB Memorandum M-25-21's directive to "redefine AI governance as an enabler of effective and safe innovation." The Department believes clear, predictable governance provides the essential "guardrails" necessary for developers and program managers to move with speed and confidence, secure in the knowledge that they are operating within safe and compliant boundaries. This framework is characterized by senior-level commitment, a lean operational model, and a clear focus on empowering the workforce. The Deputy Secretary of Transportation chairs the Department's AI Governance Board, signaling that AI is a top-tier departmental priority and ensuring that governance decisions are made with full executive authority and mission context. This senior-level engagement is critical for breaking down institutional barriers and driving enterprise-wide adoption.

The Department has embraced efficiency mandates by designating an existing, senior-level body—the Nontraditional and Emerging Transportation Technology (NETT) Council—to serve as its AI Governance Board. The NETT Council was codified in the Infrastructure Investment and Jobs Act (IIJA) as a cross-organizational body for coordinating the Department's response to nontraditional and emerging transportation technologies. Its required membership aligns closely with the intended composition of AI governance boards envisioned in M-25-21, while allowing flexibility for the Secretary to add members at their discretion. This decision reflects a commitment to lean, efficient governance that leverages existing structures rather than adding new layers of administration. While the NETT Council serves as the Department's overarching AI Governance Board, its work in this capacity is supported by task-specific sub-groups, described further below.

The governance model is federated, designed to empower trained and accountable officials at the lowest appropriate level to manage risk for most AI Use Cases, with oversight from the Department's leadership. This distributed accountability fosters a culture of ownership and agility. At the same time, the framework maintains strong, centralized oversight for high impact use cases that carry significant potential effects on public safety or rights (e.g. privacy, civil rights, civil liberties). These cases receive an appropriately higher level of scrutiny from a dedicated committee of senior experts, the SR2 Committee (described further below), before they can be deployed. This tiered approach ensures that the level of governance is always proportionate to the level of risk. A second sub-group, the AI Coordination and Activities (AICA) working group, supports the broader AI coordination function of the AI Governance Board. This structure, supported by a comprehensive set of updated agency policies, is purpose-built to slash red tape, foster clear accountability, and drive the adoption of AI to improve mission outcomes for the American people, fully aligning with the goals of M-25-21.

AI GOVERNANCE BOARD

The Department's AI governance body is the NETT Council. This existing, senior-level body was formally chartered to assume AI governance responsibilities as the DOT AI Governance Board, a strategic decision that prevented the creation of new bureaucracy and aligned with M-25-21's emphasis on avoiding "bureaucratic bottlenecks." This choice implicitly defines AI as a core component of future transportation technology, not merely as an IT issue, elevating its strategic importance.

The NETT Council was originally established in December 2018 as an internal DOT body for nontraditional and emerging transportation technologies not yet established enough to fit into modal categories, or that require new policy approaches. The NETT Council was later authorized as described in section 25008 of the Infrastructure Investment and Jobs Act (IIJA, Pub. L. No. 117-58), codified at 49 U.S.C. § 313.

Consistent with M-25-21 Section 3(a)(ii)(A), the DOT AI Governance Board is chaired by the Deputy Secretary, with the agency CDAIO serving as the vice-chair of those meetings, ensuring top-level leadership and accountability. The Board's membership includes senior representation from all key stakeholder offices as required by M-25-21, including the Chief Information Officer, General Counsel, and officials responsible for civil rights, human resources, budget, procurement, and representatives from the OAs.

As of the publication of this compliance plan, the AI Governance Board's work is supported by two task-specific sub-groups:

- AI Coordination and Activities (AICA) Working Group – Senior staff from program and research offices across the Department, assembled to coordinate the development and implementation of Departmental AI strategy and key initiatives.
- Safety, Rights, and Security Review (SR2) Committee – Senior experts in safety, security, privacy, and civil rights convened to review and approve the operational deployment of all AI Use Cases on behalf of and with oversight from the AI Governance Board. The SR2 Committee reports on its activities to the AI Governance Board at each meeting.

The Department will continue to assess and update this structure as needed. Adjustments to the structure will be reflected in future updates to this Compliance Plan.

The Board's plan to achieve its expected outcomes focuses on accelerating responsible AI adoption. It achieves this goal by reviewing and approving all high-level AI policies to ensure they promote innovation while managing risk, serving as the final arbiter for risk acceptance on the most consequential AI Use Cases, and championing barrier-removal initiatives like the AI Accelerator Roadmap. The NETT Council plans to continue its practice of consulting with external experts to broaden its perspective. This outreach is achieved through formal mechanisms

like Requests for Information (RFIs) inviting experts from research centers, academic institutions, industry, and civil society to brief the Council on emerging technologies and best practices, in accordance with applicable law.

AGENCY POLICIES

The CDAIO, operating through the AI Coordination and Activities (AICA) Working Group, is leading a comprehensive review of all existing internal AI principles, guidelines, and policies to ensure full consistency with M-25-21. This planned effort includes updating the AI Strategy, the Department's AI minimum risk management guidance, and the AI Use Case Inventory Policy.

The Department has already reviewed or is in the process of reviewing several internal policies to facilitate AI use:

- **IT Infrastructure:** The Office of the Chief Information Officer (OCIO) is reviewing its policies to accelerate Authorizations to Operate (ATOs) for the shared services AI platforms within the AI Accelerator Roadmap (ART Network, OPSLAB, and TrAIN), accelerating access for researchers and developers.
- **Data:** The Chief Data and AI Officer (CDAIO) is reviewing existing data governance “share by default” policies for application to AI deliverables.
- **Cybersecurity:** The OCIO is developing a framework to enable continuous oversight for AI applications deployed within the monitored TrAIN environment, which will improve security along with current security review processes.
- **Privacy:** The Department's Privacy Office is reviewing its Privacy Impact Assessment (PIA) process to provide any necessary specific templates and guidance for AI systems, streamlining compliance for development teams.
- **DOT** is in the final stages of developing its internal Generative AI Use Guidance, a process led by the CDAIO in collaboration with key OAs and offices. This development directly fulfills the requirement in M-25-21 Section 3(b)(iv). To ensure generative AI is used without posing undue risk, the Guidance establishes clear safeguards and oversight mechanisms.

FOSTERING PUBLIC TRUST IN FEDERAL USE OF AI

The U.S. Department of Transportation affirms its unwavering commitment to fostering public trust, a cornerstone of the M-25-21 framework and a prerequisite for the successful adoption of any transformative technology. The Department has established a clear, rigorous, and transparent process for managing risks from AI, demonstrating that our pursuit of innovation is built upon a foundation of safety and accountability including maintaining strong safeguards for civil rights, civil liberties, and privacy. This approach rejects a choice between speed and responsibility, recognizing instead that a well-defined and consistently executed risk management process is precisely what enables confident and accelerated deployment. By providing our innovators with clear rules of the road, we empower them to move forward boldly. The Department is integrating risk management into every stage of the AI lifecycle—from initial concept to retirement—through the mandatory governance checkpoints of the AI Accelerator Roadmap and the Transportation Use Case Knowledge Repository (TrUCKR). This practice ensures that risk is not an afterthought but a continuous consideration throughout development, testing, and operation.

Our framework employs a tiered approach to risk, ensuring that the level of oversight is always proportionate to the potential for impact. While all AI applications are subject to baseline oversight and tracking, those designated as "high impact" receive an appropriately higher level of scrutiny from a committee of senior experts specializing in safety, security, rights, and legal compliance. This approach ensures that our most consequential systems are vetted with the utmost rigor. The linchpin of this entire process is TrUCKR, a system that transforms abstract policies into concrete, auditable actions. TrUCKR will continue to be developed to serve as the single, authoritative system of record where high-impact determinations are made, waivers are justified and tracked, risk mitigation plans are documented as a prerequisite for deployment, and the final inventory is generated. This system creates a fully transparent and traceable governance history for every AI system. Finally, the Department has established robust controls and clear lines of authority to ensure that any AI system that fails to meet these high standards is promptly and safely removed from operation, guaranteeing that the trust of the American people is always protected.

DETERMINATION OF PRESUMED HIGH-IMPACT AI

DOT's process to determine which AI use cases are high impact is a multi-stage, documented procedure that leverages the governance structures defined in the AI Accelerator Roadmap. This process is designed to ensure a thorough review consistent with the definition of "High-Impact AI" in Section 5 of the Appendix to M-25-21.

1. **Initial Self-Assessment:** The journey begins when the sponsoring OA makes an initial assessment upon entering the use case into the TrUCKR.
2. **CDAIO Review:** The CDAIO reviews this initial assessment during the "Use Case Clearance" stage of the Roadmap. This review, which occurs before a project is granted access to the OPSLAB sandbox environment or the TrAIN development and testing environment, uses the presumed high-impact categories in M-25-21 Section 6 as its primary criteria.
3. **Safety, Rights and Security Review (SR2) Committee Consultation:** For any use case that may be high impact, the CDAIO formally consults with the SR2 Committee. This Committee provides expert advice on the potential for legal, material, or significant effects on rights or safety.
4. **Final Determination:** The CDAIO makes the final determination, which is formally recorded in TrUCKR. OAs have the right to appeal this determination to the NETT Council (the Department's AI Governance Board). The SR2 Committee will also report on its findings during each AI Governance Board meeting.

The Department's process for waivers of denial of a use case is equally rigorous. An OA can request a waiver from one or more of the minimum risk management practices by submitting a detailed justification within TrUCKR. The request is formally reviewed by the CDAIO and the SR2 Committee. The CDAIO makes the final decision, which must be based on a written determination that compliance would "create an unacceptable impediment to critical agency operations," as stipulated in M-25-21 Section 4(a)(ii). All waivers, denials, and revocations are centrally and transparently tracked in TrUCKR. The CDAIO is personally responsible for certifying the ongoing validity of each waiver on an annual basis and for tracking them, as required by M-25-21.

IMPLEMENTATION OF RISK MANAGEMENT PRACTICES AND TERMINATION OF NON-COMPLIANT AI

The Department plans to document and validate the implementation of minimum risk management practices through the mandatory, gated checkpoints of the AI Accelerator Roadmap. TrUCKR serves as the official system of record for this process. Before any high impact use case can be deployed to the TrAIN production environment, the sponsoring OA is required to complete and upload a detailed Minimum Risk Management Mitigation Plan and an AI Impact Assessment into TrUCKR. The CDAIO and the SR2 Committee must then formally review and approve these documents before granting authorization to operate in production.

This structure establishes clear lines of responsibility for implementation and oversight. The sponsoring OA is primarily responsible for implementing the risk management practices for its specific use cases. The CDAIO, with formal advice from the SR2 Committee, is responsible for independent oversight, validation, and granting

the final ATO for all high-impact systems. This structure creates a crucial separation of duties between the development function and the oversight function.

The primary control DOT has put in place to prevent non-compliant high-impact AI from being deployed is the gated structure of the AI Accelerator Roadmap. An AI system is technically and procedurally blocked from deployment into the TrAIN production environment until it receives a formal, documented clearance from the CDAIO within the TrUCKR system. This requirement creates an auditable chokepoint that prevents unauthorized or non-compliant AI from going live.

The Department's intended process to terminate non-compliant AI is built on continuous monitoring and clear authority. All high-impact AI systems in production are subject to continuous monitoring within the TrAIN environment, where automated tools will eventually report on model performance, data drift, and compliance with documented usage parameters. If this monitoring reveals that a system is out of compliance with minimum risk standards, the CDAIO, in consultation with the SR2 Committee, will direct the Chief Product and Technology Officer to suspend operations of the AI functionality immediately. As a prerequisite for deployment, all high-impact use cases are required to maintain an alternative, non-AI-dependent process. Upon suspension of the AI, the OA will immediately revert to this manual or alternative process to ensure the continuity of critical operations. The AI system will not be cleared to resume operations until the OA demonstrates that compliance has been fully restored and receives a new clearance to operate from the CDAIO. If minimum risk standards cannot be met, the system's ATO will be permanently revoked, and the system will be formally terminated and decommissioned.