Department of Housing and Urban Development (HUD) Compliance Plan for OMB Memorandum M-25-21

September 2025

Prepared by Eric Sidle, HUD Chief AI Officer

Purpose

The AI in Government Act of 20201 and Office of Management and Budget (OMB) Memorandum M-25-21, "Accelerating Federal Use of AI through Innovation, Governance, and Public Trust", direct 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 Housing and Urban Development (HUD) has developed its AI Compliance Plan, which outlines HUD'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 Innovation

Removing Barriers to the Responsible Use of AI

HUD faces various barriers to responsible use of AI and has taken steps to mitigate these barriers. Barriers and corresponding mitigation include:

Barrier	**Mitigation**
Lack of specialized expertise and skills in AI	HUD's AI Strategy includes developing an AI-skilled and ready workforce, to include both AI practitioners and users. HUD has onboarded new federal staff with AI expertise and has plans to further build out this capacity. This includes aligning AI hiring and training efforts with HUD's broader Human Capital strategy.
Limited AI tech stack	HUD's AI Strategy includes plans to develop secure and scalable infrastructure to support the full AI lifecycle – from ideation, development, and deployment to ongoing monitoring – and integrating it with HUD's overall IT ecosystem, enterprise architecture, data governance, and security and privacy pipeline. HUD has also been actively engaging with other federal agencies to identify potential partnerships, to include leveraging government-wide AI services and sharing best practices. HUD is utilizing our governance processes to introduce appropriate tools to rapidly develop, test, and maintain AI applications.

Existing acquisition processes not fully tailored to unique requirements and risks of AI procurement	HUD has convened a cross-functional team dedicated to streamlining the acquisition process in alignment with M-25-22, "Driving Efficient Acquisition of Artificial Intelligence in Government". The team is identifying and inserting best practices and tools into the existing HUD acquisition process, such as how market research and source selections are conducted. The team is also developing contract terms that incorporate considerations, such as data and IP rights, model performance, and responsible AI use. HUD also plans to partner with other government agencies to leverage available centralized AI tool procurement standards and vehicles.

Sharing and Reuse

HUD is enhancing and streamlining its overall processes to expedite and improve value delivery, including business need intake and technical solutioning. All new business requests come through the Office of the Chief Information Officer (OCIO), which determines the appropriate technical solution. When evaluating solutions, the OCIO first considers existing technical capabilities – to include AI code and models – and reuses them where appropriate. If no fit exists, the OCIO will develop a new solution, which is then added to HUD's common source code repository for future use.

To further promote sharing and reuse, HUD is also building out its Continuous Integration and Continuous Delivery/Development (CI/CD) and Machine Learning Operations (MLOps) pipelines that integrate with the common source code repository.

Additionally, HUD through its Chief Data Officer has recently kicked off the effort to build out an enterprise data inventory, as well as an enterprise data strategy and governance. These efforts will expand data sharing and provide the underlying AI-safe and AI-ready data needed to enable broader reuse of AI solutions.

AI Talent

As part of HUD's AI strategy, HUD is focused on the development of an AI-skilled workforce to support the deployment and oversight of AI at HUD. The plan includes:

- Defining AI workforce needs: Identify critical AI roles and mapping these to mission areas that will have optimal impact.
- Recruiting and hiring AI talent: Develop AI-specific position descriptions and partner with various groups (e.g., government entities, including state and local, academia, and professional organizations) to attract AI talent.
- Training and upskilling the HUD workforce: Develop training plans for technical staff, leveraging available training programs and resources and supplementing where needed.
- Retaining AI talent: Develop and offer flexible career paths and growth opportunities through details and new assignments, and empower and engage AI staff.

As HUD builds out its AI talent, it will be paying particular attention to obtaining expertise in areas such as building Machine Learning Operations (MLOps) into the overall Continuous Integration/Continuous Delivery (CI/CD) pipeline, model development, evaluation, and testing, monitoring and risk management

(e.g., tracking model drift, bias, and performance degradation), and security and compliance (e.g., related to data privacy and regulatory compliance). These skillsets will be useful and highly impactful for AI implementations across all of HUD's program areas.

Improving AI Governance

AI Governance Board

In compliance with M-25-21, HUD reconvened its AI governance body in July 2025. To streamline the process, HUD leveraged its existing Principles Investment Review Board (PIRB) and expanded its responsibilities to include governance of AI at HUD. The PIRB is chaired by the HUD Deputy Secretary and the Chief Information Officer (CIO) who is also the Chief Artificial Intelligence Officer (CAIO). The PIRB membership is comprised of principles from each of the HUD program offices, including from the Office of Housing, Community Planning and Development, Public and Indian Housing, Field Policy and Management, Policy Development and Research, Fair Housing and Equal Opportunity, Ginnie Mae, as well as finance, IT, legal counsel, privacy, procurement, and human capital.

To achieve its expected outcomes, including compliance with executive orders and OMB memos, while staying nimble and streamlined, the AI governance board has and will continue to establish targeted working groups to address specific needs as they arise. For instance, the board stood up a cross-functional AI Strategy working group with the objective of developing the agency's AI Strategy.

In addition, the HUD CAIO is also a member of the Federal CAIO Council. The CAIO has worked and will continue to engage with other federal agencies to share knowledge and ideas and find partnership opportunities.

Agency Policies

HUD is updating its internal AI guidelines and processes, which are consistent with M-25-21 and M-25-22. This includes revamping HUD's underlying technology infrastructure and processes, such as revamping the CI/CD pipeline to include MLOps. HUD is also improving and streamlining its overall system development lifecycle, which includes improved business intake and technical solutioning, as well as integration with the new AI high-impact review process and the streamlined AI acquisition process aligned with OMB M-25-22. HUD also recognizes the importance of a strong data foundation and is developing new data principles that include ensuring data is AI-ready and AI-safe to enable responsible AI innovation, while safeguarding HUD's information assets. In addition, as HUD runs its new AI use cases through its processes, HUD will be responsive to lessons learned and insights and iteratively revise processes and guidelines as needed.

HUD currently provides a disclaimer to HUD staff when they attempt to access AI tools to ensure users do not input sensitive information and to validate the outputs of the AI system. HUD will be working on broader generative AI guidance.

AI Use Case Inventory

HUD has a process for collecting AI uses cases, which is integrated into HUD's overall IT governance process as follows:

1. Requestor submits business need
2. OCIO reviews the business need, including an Enterprise Architecture (EA) review. As part of the solutioning, the team will determine whether the solution will involve non-embedded AI.
3. If so, the proposed technical solution will go through the AI high-impact review.
4. The results of the high-impact assessment will be presented to HUD's Technical Review Committee (TRC) for a technical review and decision.
5. If the solution is approved, then it will be added to HUD's AI Use Case Inventory.
6. If the use case is considered high-impact, the system owner will be required to follow specified risk management requirements. These outputs will be reviewed during the appropriate gate review(s) in the system development lifecycle process. The AI Use Case Inventory will reflect these updates.

HUD's AI Use Case Inventory is tracked in a centralized document repository. The AI Use Case Inventory will also be validated with the respective owners prior to posting publicly.

Fostering Public Trust in Federal Use of AI

Determinations of Presumed High-Impact AI

HUD has developed processes for determining which AI use cases are high-impact and for determining if one or more of the minimum risk management practices can be waived. The processes are as follows:

<u>High Impact Determination Process</u>

1. AI system owner completes the AI Risk Review checklist. Questions include:
   - Description of the use case and problem it solves
   - Data / inputs of the AI to make decisions or recommendations
   - Outputs of the AI and how the outputs are used
   - Impacts, including affects on eligibility, access, benefits, rights, and safety
2. AI risk team reviews the responses with the AI Use Case Owner and if any of the listed impacts are 'Yes', then the AI use case is deemed high-impact.
3. If the AI use case is high-impact, the system owner is required to meet a minimum set of risk management requirements, which are consistent with M-25-21.
4. The system owner can request to waive one or more of the risk management requirements (see below).

Risk Management Waiver Process

As part of the AI high-impact determination process, the AI use case owner has the option to submit a request to waive one or more of the risk mitigation requirements for a limited time per the following process:

1. AI use case owner submits the AI risk management requirement waiver form, which includes details on the waiver request (e.g., requested waiver, justification, duration of waiver, proposed alternative controls, and plan for meeting the risk management requirements by the end of the waiver period)
2. AI risk team reviews the waiver request, and has a set of questions to ask.
    - Example: If the waiver is requesting to waive the Pre-deployment testing requirement, the AI risk team will ask questions to include: "Are there alternative controls in place to manage risks during the waiver period?", "Can partial testing be done to reduce risk?", "Is there a plan and timeline to complete testing?". The waiver option for this would be "Waiver can only be considered if there is: A plan and timeline to complete testing, **and** Risk mitigation in place until testing is completed".
3. AI risk team provides its recommendation.
    a. If the AI risk team does not approve the waiver request, the AI use case owner is notified and they will need to adhere to the risk mitigation practices as defined.
    b. If the waiver request is approved by the AI risk team, they will update the waiver determination form with their recommendation to the CAIO.
4. The CAIO reviews the request and recommendation and makes their final determination.
5. All waivers, with decisions, are tracked and stored centrally.

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

As part of its effort to streamline and simplify its processes, HUD is incorporating the risk management requirements into the overall system development lifecycle process. Specifically, if a system or component is identified as high-impact, its risk management requirements compliance will be validated as part of the appropriate system development lifecycle gate review(s). This includes a gate review prior to deployment, so if a system does not meet the minimum risk management requirements, it will not be authorized to be deployed. HUD is also planning to update its IT security and privacy assessment processes – including the Authority to Operate (ATO) process – to incorporate high-impact considerations. The plan will be to review a high-impact AI system or component as part of the regular ATO review and re-certification process and, if the AI system or component is non-compliant, will not receive its re-certification and will need to complete a plan of action and milestones.