Tennessee Valley Authority (TVA) Compliance Plan for OMB Memorandum M-25-21 – July 2025

- Prepared by KC Carnes, Vice President, Cybersecurity & Chief Information Security Officer

1. Driving AI Innovation

Removing Barriers to the Responsible Use of AI

TVA has identified potential barriers to the responsible use of AI and is currently taking proactive action to remedy and mitigate.

TVA faces a challenge regarding the limitations of current technology for developing and scaling AI projects. Due to AI projects requiring significant computing and storage resources to train and run models effectively, TVA is creating robust and scalable infrastructure to access high quality, governed data and implement AI projects through cloud-based solutioning. Additionally, through enabling Machine Learning Operations (MLOps) capabilities within this cloud-based infrastructure, TVA aims to centralize and streamline the development, technology, and operations of AI and machine learning projects while promoting and enabling rapid experimentation and efficient, scalable innovation.

Model explainability is a crucial yet challenging aspect of the development process for establishing trust in AI systems. TVA supports the implementation of technology and standards that will provide interpretability, what-if and stress testing, bias detection, error detection, and system output validation so that users have insights into the factors influencing a decision, prediction, or model output leading to increased trust and acceptance of AI technology.

After developing machine learning models, TVA faces challenges in deploying and monitoring them across multiple environments, which can hinder risk mitigation and the effective management of AI performance. By implementing MLOps monitoring capabilities, TVA aims to foster a robust environment for the automated monitoring and evaluation of the MLOps pipeline and deployment of AI applications. This includes monitoring of model performance, quality, bias, drift, usage, and the detection of any issues or anomalies to ensure ongoing accuracy, security, and reliability.

Furthermore, data privacy and security concerns are critical barriers. Ensuring that these concerns are addressed effectively and efficiently in all instances of various types of AI can be challenging. TVA is approaching these challenges with different initiatives such as AI governance, policy frameworks, and technical guidance that drive the stringent data governance policies to include tactical measures that align with technology security requirements and guidelines in other Executive Orders and TVA policies.

Addressing ethical concerns and mitigating biases in AI models is essential for the responsible use of AI. To tackle this challenge, TVA is establishing internal frameworks to support the oversight of ethical concerns to include an AI ethics program. The intent of the AI ethics program is oversight of AI to ensure that ethical considerations for AI projects are thoroughly reviewed and align with TVA's standards. Bias detection and mitigation

mechanisms will be integrated into the development life cycle for AI technologies to support the AI ethics program from a technical perspective.

Addressing the challenges and barriers for responsible use of AI at TVA will ensure its successful and safe use from design to deployment. TVA is aware that these challenges will change and is taking action to ensure awareness of the dynamic nature of AI and ensuring alignment with regulatory requirements while creating an environment for innovation and operational advancement.

Generative AI (GenAI) technologies pose new challenges and risks for TVA, including the ideology of hallucinations or invented and inaccurate responses that may be biased or based on biased training data. Additionally, if sensitive data is provided and stored for further use and training by GenAI, the technology poses a threat to the security and privacy of TVA data. TVA is actively exploring GenAI use cases and methodologies that ensure the safe, secure, and trustworthy use of the technology.

Sharing and Reuse

TVA is committed to ensuring that custom-developed AI code, including models and model artifacts, is shared per regulatory requirements. This process is governed through a structured and centralized approach facilitated by the use of an AI governance platform.

TVA's selected AI governance platform serves as a central place for managing all data and analytic initiatives pertaining to AI development. This includes overseeing the life cycle of the AI projects, their respective models, versions, data sets, and documentation. Through centralizing these governance activities, TVA ensures that all AI developments are consistently tracked, reproducible, and managed in alignment with defined policy and development requirements. TVA will control the deployment of projects with a built-in sign-off process, ensuring authorization and validation prior to deployment. This will further enable TVA to effectively monitor deployed models and analyze their performance.

TVA encourages and incentivizes the transparency and sharing of AI code, models, and data with the public, where permitted and feasible, according to applicable regulations and laws to foster innovation and collaboration with external entities and the public. Several initiatives and business offices within TVA are responsible for the coordination of these efforts.

TVA will investigate the development of an incentive program, such as recognition awards and professional development opportunities, to encourage TVA employees to share their AI developments. Additionally, TVA will consider the best path forward for public collaboration platforms like GitHub and open data portals to facilitate external engagement. Community workshops will also promote best practices and showcase successful projects. TVA's Enterprise Analytics and AI team (EA&AI), along with the support of various business partners, ensures technical and policy support for these initiatives, aligning them with regulatory requirements and organizational goals.

AI Talent

TVA is actively pursuing initiatives to increase AI talent, recognizing the critical importance of skilled personnel in driving AI innovation and responsible use. To achieve this, TVA will leverage specific hiring authorities, establish AI-focused teams, and target key skillsets and expertise levels. To consider all potential impacts to AI adoption and adaptation, TVA will conduct workforce assessments to address the potential impacts of AI automations and augmentation across job functions and roles.

TVA's EA&AI team currently serves to empower business units with data and AI knowledge, skills, and expertise. They ensure predictable and repeatable usage of data and AI services across the enterprise, supporting innovation and collaboration. Operating at an enterprise level, providing centralized support for promoting and developing TVA's data and AI capabilities, the EA&AI team will improve business outcomes.

The EA&AI team supports AI development and facilitates knowledge sharing across the enterprise through business unit outreach and collaboration. The outreach covers AI technologies to educate users and business entities with baseline AI knowledge that can help drive forward-thinking and innovation for initiative advancement. They also play a central role in creating a community of practitioners to enable the cultural change outlined in TVA's strategy. Moreover, they encourage responsible AI use of self-service solutions and provide expert consultation for coaching and guidance on business process needs.

To meet the fast-growing and changing AI landscape, TVA will develop specialized AI learning opportunities to meet the diverse needs of its workforce. Training will enhance data and AI literacy, educate on AI concepts and responsible AI development, and ensure users can adequately mitigate risks associated with AI tools. As AI is integrated into TVA's infrastructure and technology platforms, it is integral that the workforce understands AI concepts, methods, and technologies while knowing how to safely use and apply them. TVA will continue to increase AI fluency for both technical and non-technical users by enhancing its AI curriculum training pathways. This includes expanding its current training to include Data Fluency, AI Fluency, and AI Ethic programs to ensure responsible understanding, development, and use of AI tools and applications.

TVA will strive to cultivate a culture of innovation and forward-thinking to attract and retain top AI talent. Environments that offer opportunities for growth, creativity, and learning are highly appealing to talented professionals. By fostering such a culture, TVA aims to become a desirable destination for top AI professionals.

TVA aims to establish an AI Talent initiative to oversee AI talent and ensure alignment with organizational goals. This initiative will be responsible for coordinating recruitment efforts, managing training and development programs, and fostering a culture of continuous learning and innovation within the AI teams. TVA aims to build a robust and skilled AI workforce capable of advancing the organization's AI capabilities, ensuring responsible use of AI technologies, and maintaining a competitive edge in the rapidly evolving field of artificial intelligence.

2. Improving AI Governance

AI Governance Board

TVA maintains a governance committee that was established as a subordinate committee of TVA's Enterprise Risk Council (ERC) to provide comprehensive risk oversight of TVA's security, artificial intelligence, privacy, and technology risks consistent with TVA's mission, strategic imperatives, and approved financial and operational plans.

The governance committee is responsible and accountable for the evaluation, approval, and monitoring of risk management processes and infrastructure, in accordance with the ERC's strategic direction, for TVA's security, artificial intelligence, privacy, and technology risks to protect TVA's interests. The committee oversees TVA's strategies and risk mitigation measures to ensure cross-functional alignment and consistency with established laws, regulations, guidelines, and best practices.

TVA's governance committee includes members from all critical business components, including agency officials from Financial Services, TVA Nuclear, Chief Operating Office, and Human Resources.

The expectations for the governance committee are to review and approve or reject proposed risk responses in alignment with TVA Policy for Enterprise Risk Management, when exposure could pose a significant risk impact to TVA. The governance committee supports both the consultation efforts with external entities on AI and across the Federal Government by selecting appropriate delegates to participate in relevant interagency communities.

Agency Policies

TVA published an agency-wide policy and technical standards to set guardrails and requirements that align with federal requirements and best practices. These policies and technical standards were written in favor of forward leaning innovation, operational availability, and mission enabling. TVA's published policy is used to maximize investment opportunities that enable AI adoption, while ensuring a balance of safety and security measures. Governance of AI has effectively been designed to ensure risk management is embedded in all AI use cases without unintentionally hindering or delaying the ability to reach a return on investment in the technology.

Multiple business units within the agency support the development of creating and disseminating appropriate documentation on AI best practices, from design to deployment and use. The various AI documentation is published within centralized and standardized documentation repositories for TVA personnel to access guidelines, standard requirements, and organizational policy and strategies. The resources are maintained and updated as needed to dynamically support the responsible use of AI.

In addition, TVA's EA&AI serves to empower the business with digital AI technologies and capabilities while providing knowledge management and centralized information. The EA&AI's collaborative spaces (such as repositories and discussion forums) help to answer questions, provide documentation and guidance on best practices, retain key lessons, and share experiences with peers and experts in order to support responsible innovation and collaboration.

Within the EA&AI, TVA has formed an AI Enablement Working Group for

communication and knowledge exchange with the goal of understanding how the enterprise can responsibly leverage AI to solve complex business challenges. The cross-functional group serves as a forum for sharing and understanding risks and regulations, emerging trends, AI initiatives, AI technologies, and identifying opportunities for AI integration within various business functions. By providing a focal point for knowledge sharing, the EA&AI seeks to improve the way TVA leverages technology to solve problems and improve operations.

AI Use Case Inventory

TVA has established procedures for collecting AI use cases at the beginning of their life cycle. This involves engagement and outreach sessions with business partners across TVA to gather feedback on business problems, documentation of the AI types, design, prioritizations, and affiliated technology. These are then added to TVA's standard asset management platforms for continuous trackability. Updates are made to these assets as new information becomes available during normal change management processes. The EA&AI team educates and encourages partners to consider AI technologies for solving problems and advancing opportunities.

EA&AI conducts value and feasibility studies to assess the practical application of potential AI use cases, determines the type of AI, design of the system, complexity of implementation, and priority in alignment with other business initiatives. The team documents existing and planned AI initiatives, detailing their objectives, methodologies, and anticipated outcomes. Regular update sessions are scheduled to document progress and modifications to AI initiatives, ensuring accuracy of AI initiative tracking. TVA maintains asset tracking of AI technologies in alignment with the overall strategy for asset management, supported by MLOps and AI governance platform for enhanced tracking and visibility of active deployments.

After EA&AI initially reviews an AI use case, TVA's Cybersecurity team will evaluate the proposed AI use case to ensure compliance with federal directives, such as executive orders and memorandums. During this review, TVA assesses whether the AI technology is considered high impact and if the minimum risk practices are required. An annual evaluation is conducted for existing AI technologies to ensure there are no changes with the intended use that alter the tools impact. This review is considered a risk review and includes reporting recommendations and analysis of the AI technology to the AI governance committee.

The framework for evaluating the potential risks associated with the AI technologies ensures alignment with strategic goals and regulatory frameworks, adherence to ethical guidelines, privacy considerations, and operational efficiency. By continuously conducting these reviews, TVA maintains the integrity and validity of its AI inventory and the necessity of AI use within the agency.

3. Fostering Public Trust in Federal Use of AI

Determinations of Presumed High-Impact AI

TVA has a structured evaluation process that effectively determines in the early stages of an AI use case request, if the solution will be High-Impact. TVA leverages the criteria defined within OMB M-25-21 for High-Impact AI. The EA&AI team will support the initial evaluation process through determination of the impact and feasibility of the use case request. Should the use case be identified to move forward, TVA Cybersecurity will officially determine the High-Impact status and work to ensure the minimum risk practices are appropriately implemented with the use case.

This process workflow will be handled within TVA's Information Technology Service Management System, ensuring that the request is routed to appropriate technology groups and those involved are clearly identified. Requestors can then carefully track the status and conclusion of the evaluation process. The outcome of Cybersecurity's comprehensive review process will determine next steps.

Denial: Cybersecurity has determined that the use-case should not move forward. No further action is required.

Approval: Cybersecurity has determined that the use-case is approved to move forward into design phases.

Risk Evaluation Required: Due to elevated risk and/or High-Impact status, Cybersecurity has determined that the use-case requires elevated analysis and final approval from the AI governance committee prior to moving into the design phase.

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

TVA ensures the documentation and validation of minimum risk management practices through both manual and automated processes throughout the AI technology life cycle. Technical standards define requirements for designing, developing, and deploying AI, specifying roles, processing requirements, and responsibilities. These standards provide both tactical and high-level guidance within the MLOps framework for AI governance and monitoring.

Before any implementation of an AI use case application, system, or technology, TVA conducts adequate testing to ensure the AI, as well as components that rely on it, will work in its intended real-world context. Through appropriate testing, TVA can demonstrate that the AI system will achieve its expected benefits and that associated risks will be sufficiently mitigated, or else it will be determined the AI system should not be used. Prior to operationalizing AI systems, guardrails and monitors are implemented where any risk issues may arise (e.g., quality monitoring, bias monitoring, performance monitoring, usage monitoring, etc., with thresholds to alert users). After deployment, training, safeguards, and ongoing monitoring are implemented to evaluate AI functionality degradation and to detect changes in its overall impact status. TVA will introduce AI observability techniques to oversee GenAI technologies that promote understanding and monitoring of model output.

Continuous validation and reviews ensure adherence to risk management practices, with formal reviews occurring annually and more frequently, as needed, based on specific risk

management strategies. TVA's AI policy identifies business units responsible for implementing and overseeing these practices, ensuring alignment with federal, regulatory, and internal risk management requirements.

All of TVA's technology environments are managed and operated in alignment with the overall risk strategy. Specific to AI and management of risk, TVA is currently implementing automated and manual measures that will detect if an AI technology is non-compliant and High-Impact AI. TVA will accomplish the determination of non-compliant AI technologies through continuous reviews, integrated governance blueprints and frameworks through the development process, automated compliance checks, and regular monitoring and analysis. TVA's use of its AI governance platform integrated with MLOps will support efforts in proactively detecting policy violations and ensure minimum risk practices are applied to High-Impact AI.

TVA closely models the Risk Management Framework for AI within the agency. This ensures that the necessary policies, processes, and procedures are in place across the agency, enabling TVA to map controls, measure and manage AI risks, and validate the effectiveness of the AI controls through auditability and transparency in monitoring practices. Should initial, out-of-band, or annual reviews determine that a specific AI technology is non-compliant with regulatory requirements for High-Impact AI, TVA will promptly initiate and prioritize a Plan of Action and Milestones (POA&M) to ensure the technology is effectively and timely brought back into compliance or is suspended for as long as necessary to comply with appropriate regulatory requirements.