TITLE: Tennessee Valley Authority Compliance Plan for OMB Memoranda M-24-10 – September 2024

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

1. STRENGTHENING AI GOVERNANCE

General

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

The Tennessee Valley Authority (TVA) is proactively taking action to remain aligned with requirements and objectives as defined within the Office of Management and Budget (OMB) M-24-10, industry standards, and best practices for the secure design, development, and deployment of Artificial Intelligence (AI) technologies. Specifically, TVA has published an enterprise-wide policy that defines the roles and responsibilities for those who use AI at TVA. The policy identifies the Chief AI Officer (CAIO) and Security & Technology Control Board (STCB) which will oversee the use of AI at TVA, as well as many other applicable roles that support the advancement of AI safely and ethically. TVA's efforts are centered on integrating principles of transparency, accountability, and equity defined within M-24-10 into existing and new development frameworks that ensure fairness, privacy protection, and ethical standards are upheld.

The development of this policy and the drafting of future technical standards have involved consultations and collaboration with stakeholders across many different business units at TVA. This collaboration is designed to capture diverse perspectives, insights, and ensure comprehensive approaches that meet regulatory requirements and operational needs. The focus on progressing AI forward in a safe and responsible way has led TVA to enhance policy frameworks that reflect the evolving landscape of the AI technologies and regulations. TVA will continue to adopt appropriate policy and governance practices that promote the safe and responsible use, while clearly defining procedures for AI design, development, and deployment. TVA's technical standards and future policies define the framework for governance and oversight and robust mechanisms for monitoring and evaluation. TVA will foster a culture of responsible use within the agency where innovation is balanced with ethical standards and federal requirements.

AI Governance Bodies

- Identify the offices that are represented on your agency's AI governance body.
- Describe the expected outcomes for the AI governance body and your agency's plan to achieve them.
- Describe how, if at all, your agency's AI governance body plans to consult with external experts as appropriate and consistent with applicable law. External experts are characterized as individuals outside your agency, which may include

individuals from other agencies, federally funded research and development centers, academic institutions, think tanks, industry, civil society, or labor unions.

The STCB will act as TVA's governance body, consisting of internal TVA stakeholders from Technology & Innovation, Nuclear, Treasury & Risk, Transmission & Power Supply, Power Operations, Human Resources, and TVA Police & Emergency Management. This group is responsible and accountable for the evaluation, approval, and monitoring of risk management processes and infrastructure. The STCB operates in accordance with TVA's Enterprise Risk Committee's strategic direction for TVA's security, artificial intelligence, privacy, and technology risks to protect TVA's interests.

The expectations for TVA's STCB are to review and approve proposed risk responses, in alignment with TVA Policy for Enterprise Risk Management, when a risk exposure could pose a significant risk impact to TVA. The STCB is poised to achieve several critical outcomes aimed at ensuring the responsible and effective use of AI technologies.

First, by enhancing the transparency and accountability of TVA's AI initiatives and by centralizing oversight and decision-making within the STCB, TVA can ensure that AI technologies adhere to TVA's policy and standards, extended ethical standards, and regulatory requirements.

Second, the STCB will facilitate consistency and coherence across the agency within relevant AI initiatives and projects. The STCB will provide a structured framework for assessing and managing the risks associated with AI deployments to minimize potential negative impacts on stakeholders and ensure operational availability. Interdisciplinary dialogue and best practice dissemination can harness the full potential for innovation and use of AI technology services and applications, ultimately improving TVA's efficiencies and service to the people of the Tennessee Valley.

Finally, TVA is committed to leveraging external expertise to inform and enhance decision-making processes in a manner consistent with applicable laws and regulations. TVA will consult with a diverse range of external experts, including those from other federal agencies, federally-funded research and development centers, academic institutions, industry leaders, civil society organizations, and labor unions, as needed. The efforts will be to gather insights on emerging AI trends, ethical considerations, best practices, and potential societal impacts. This approach will strengthen the integrity and effectiveness of TVA's AI governance framework, promoting responsible and beneficial deployment of AI technologies across TVA operations.

AI Use Case Inventories

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

TVA is a smaller FCEB Agency and the agency process is documented within this AI Compliance Plan.

Reporting on AI Use Cases Not Subject to Inventory

- Describe your agency's process for soliciting and collecting AI use cases that meet the criteria for exclusion from being individually inventoried, as required by Section 3(a)(v) of M-24-10. *In particular, explain the process by which your agency determines whether a use case should be excluded from being individually inventoried and the criteria involved for such a determination.*
- Identify how your agency plans to periodically revisit and validate these use cases. *In particular, describe the criteria that your agency intends to use to determine whether an AI use case that previously met the exclusion criteria for individual inventorying should subsequently be added to the agency's public inventory.*

TVA is establishing 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 into 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 Enterprise Analytics & Innovation (EA&I) team educates and encourages partners to consider AI technologies for solving problems and advancing opportunities.

EA&I 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 a Machine Learning Operations (MLOps) and AI governance platform for enhanced tracking and visibility of active deployments.

[4] Examples of updates to existing use cases include moving to a different phase of the system development life cycle (e.g., development to operations) or updating the documentation for risk management activities when a rights- impacting AI or safety-impacting AI enters into use.

After EA&I 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 impacts safety or human rights using criteria defined in M-24-10. This evaluation is conducted annually for existing AI technologies to ensure there are no changes in impact. This review is considered a risk review and includes reporting recommendations and analysis of the AI technology to the STCB.

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.

During the continuous evaluation process, TVA will follow specific criteria to determine whether an AI use case that previously met exclusion criteria for individual inventorying should be added and reported on as described within M-24-10. Attributes that may justify the necessity for individual inventorying include:

Change to impact: The technology or purpose has changed and now effectuates one or more of the criteria for safety- or rights-impacting, as defined within M-24-10.

Regulatory and compliance updates: There have been updates to the applicable memorandums or executive order requirements that trigger inclusion requirements for the technology.

Operational changes: The technology scope has changed, warranting an out-of-band assessment of the technology.

Stakeholder feedback: Internal agency stakeholders, such as the STCB, have reviewed the technology and its use, determining that it should be reported on based on public interest or business initiatives and relevant objectives.

Ethical and privacy concerns: Should ethical or privacy concerns arise from the use of the technology, TVA may take action to include in public inventory for transparency and applicability to regulatory requirements.

2. ADVANCING RESPONSIBLE AI INNOVATION

Removing Barriers to the Responsible Use of AI

- Describe any barriers to the responsible use of AI that your agency has identified, as well as any steps your agency has taken (or plans to take) to mitigate or remove these identified barriers.[6] *In particular, elaborate on whether your agency is*

[6] Section 4(b) of M-24-10 references barriers such as access and support for IT infrastructure, data, and cybersecurity.

addressing access to the necessary software tools, open-source libraries, and deployment and monitoring capabilities to rapidly develop, test, and maintain AI applications.

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

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

TVA faces a challenge regarding the limitations of current technology for developing and scaling AI projects. Because AI projects require 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 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 is supporting the implementation of technology and standards that will provide interpretability, what-if and stress testing, bias detection, error detection, and that can validate a system's output 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 EOs 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 barriers 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.

AI Talent

- Describe any planned or in-progress initiatives from your agency to increase AI talent. *In particular, reference any hiring authorities that your agency is leveraging, describe any AI- focused teams that your agency is establishing or expanding, and identify the skillsets or skill- levels that your agency is looking to attract. If your agency has designated an AI Talent Lead, identify which office they are assigned to.*
- If applicable, describe your agency's plans to provide any resources or training to develop AI talent internally and increase AI training opportunities for Federal employees. *In particular, reference any role-based AI training tracks that your agency is interested in, or actively working to develop (e.g., focusing on leadership, acquisition workforce, hiring teams, software engineers, administrative personnel or others).*

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, establishing focused AI teams, and targeting key skillsets and expertise levels. To consider all potential impact to AI adoption and adaptation, TVA plans to conduct workforce assessments to address the potential impacts of AI automations and augmentation across job functions and roles.

TVA's Data and Analytics Center of Excellence (D&A CoE) currently serves to empower business units with data and AI knowledge, skills, and expertise. The D&A CoE ensures predictable and repeatable usage of data and AI services across the enterprise, supporting innovation and collaboration. It operates at an enterprise level, providing centralized support for promoting and developing TVA's data and AI capabilities to improve business outcomes.

The D&A CoE 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. The CoE also plays a central role in creating a community of practitioners to enable the cultural change outlined in TVA's strategy. Moreover, it encourages responsible AI use of self-service solutions and provides 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, and knows 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 a Data Fluency, AI Fluency, and AI Ethics program to ensure the 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 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.

AI Sharing and Collaboration

- Describe your agency's process for ensuring that custom-developed AI code—including models and model weights—for AI applications in active use is shared consistent with Section 4(d) of M-24-10.
- Elaborate on your agency's efforts to encourage or incentivize the sharing of code, models, and data with the public. Include a description of the relevant offices that are responsible for coordinating this work.

TVA is committed to ensuring that custom-developed AI code, including models and model artifacts, for AI applications in active use is shared according to Section 4(d) of M-24-10. 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 be able to control the deployment of projects with a built-in sign-off process, ensuring authorizations 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 offices within TVA are responsible for the coordination of these efforts.

TVA values continuous improvement as a core competency and employees, driven by the D&A CoE, will be encouraged to find opportunities to highlight and share AI developments and look for ways to scale or expand the usage of AI solutions across the enterprise. Additionally, TVA will consider the best path forward for public collaboration platforms like GitHub and open data portals to facilitate external engagement, while community workshops will promote best practices and showcase successful projects. The D&A CoE, along with the support of various business partners, ensures technical and policy support for these initiatives, aligning them with regulatory requirements and organizational goals.

Harmonization of Artificial Intelligence Requirements

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

TVA is diligently developing policy and technical standards to define and document best practices as it pertains to the innovation of AI and its governance and risk management. TVA has published enterprise-wide policy for the responsible use and governance of AI within the agency. 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 D&A CoE serves to empower the business with digital AI technologies and capabilities while providing knowledge management and centralized information. The D&A CoE'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 D&A CoE, 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 D&A CoE seeks to improve the way TVA leverages technology to solve problems and improve operations.

3. MANAGING RISKS FROM THE USE OF ARTIFICIAL INTELLIGENCE

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

- Explain the process by which your agency determines which AI use cases are rights- impacting or safety-impacting. *In particular, describe how your agency is reviewing or planning to review each current and planned use of AI to assess whether it matches the definition of safety-impacting AI or rights-impacting AI, as defined in Section 6 of M-24-10. Identify whether your agency has created additional criteria for when an AI use is safety-impacting or rights-impacting and describe such supplementary criteria.*
- If your agency has developed its own distinct criteria to guide a decision to waive one or more of the minimum risk management practices for a particular use case, describe the criteria.
- Describe your agency's process for issuing, denying, revoking, tracking, and certifying waivers for one or more of the minimum risk management practices.

TVA has a structured evaluation process that will replicate the evaluation criteria as defined within M-24-10. TVA's EA&I 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 into design stages, TVA's EA&I and Cybersecurity teams will review the AI use case to determine if it is rights- and/or safety-impacting.

This process workflow will be handled within TVA's Information Technology Service Management System, ensuring that the request is routed to appropriate 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: Cybersecurity has determined that the use-case should require STCB analysis and final approval to move forward to design.

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

- Elaborate on the controls your agency has put in place to prevent non-compliant safety- impacting or rights-impacting AI from being deployed to the public.
- Describe your agency's intended process to terminate, and effectuate that termination of, any non-compliant AI.

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 safety- or rights-impacting. TVA will accomplish the determination of non-compliant AI technologies through continuous reviews, integrated governance blueprints and framework 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 that could implicate safety- or rights-impacting AI models.

TVA closely models the Risk Management Framework for AI for use of 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 safety- and rights-impacting AI, TVA will promptly initiate and prioritize a Plan of Action and Milestones (POA&M) to ensure the technology is effectively brought back into compliance or is suspended for as long as necessary to comply with appropriate regulatory requirements.

Minimum Risk Management Practices

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

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 impact on rights and safety. 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.