Best Practices for Procuring Generative AI in Government (State DOTs and Beyond)

Requested by
Dara Wheeler, Acting Chief Data and Artificial Intelligence Officer
Author
Mehdi Moeinaddini, Senior Transportation Planner, Office of Planning, Policy, and Program Development, Division of Research, Innovation, and System Information

Jun 20, 2025

Table of Contents

Executive Summary

This review involves selecting relevant studies, news articles, and reports, and synthesizing their key themes and findings to provide an overview of best practices, opportunities, and considerations for the procurement and integration of Generative AI (GenAI) technologies within Caltrans.

Summary of Key Findings

Policy Frameworks

- GenAI procurement integrates traditional IT procurement processes with additional considerations unique to AI, such as managing algorithmic bias, security, transparency, and accountability.
- Federal guidelines from U.S. Office of Management and Budget (OMB) and General Services Administration (GSA) emphasize cross-functional governance, early involvement of privacy/security officials, and outcomes-based procurement practices.
- State governments (e.g., California, Washington) are actively developing frameworks and dedicated resources for GenAI, guided by the National Institute of Standards and Technology's (NIST) AI Risk Management Framework.

Vendor Selection and Contracting

- Vendor selection now includes AI-specific criteria like algorithmic risk assessments, fairness, bias mitigation, transparency, explainability, data quality, and security compliance.
- Vendors must demonstrate robust risk mitigation practices, and procurement teams frequently employ sandbox tests or pilots to validate vendor claims.
- Contract terms often focus on transparency, security standards (e.g., FedRAMP), data ownership and usage restrictions, and exit strategies to avoid vendor lock-in.

Data Privacy and Security

- GenAI solutions must strictly adhere to privacy laws, with contract clauses explicitly prohibiting unauthorized use of government data for model training.
- Contracts include rigorous cybersecurity standards, data management plans, and legal safeguards ensuring government rights over AI-generated outputs.
- Agencies prioritize secure environments, often utilizing government-specific cloud solutions and secure data sandboxes for GenAI experimentation.

Implementation and Integration Strategies

- Agencies typically start with clearly defined business needs and small-scale pilots to test GenAI solutions effectively.
- Effective implementation involves forming interdisciplinary teams, ensuring governance oversight, and proactively managing ethical and equity considerations.
- Training, change management, phased deployment, and careful cost control are crucial to successful integration and scaling of GenAI tools.

Monitoring and Continuous Improvement

- Continuous monitoring, regular revalidation of performance, and ongoing vendor accountability are emphasized.
- Agencies document and share experiences and best practices, promoting collective learning across the public sector through established communities of practice.

State DOTs and other agencies recognize that traditional procurement cycles cannot keep pace with rapidly evolving AI technology. Emerging practices include agile procurement, challenge-based solicitations (such as California's Request for Innovative Ideas, RFI²), modular and phased contracts, and leveraging flexible vendor pools and cooperative purchasing agreements. These methods significantly reduce procurement timelines, enhance innovation, and improve adaptability.

Implementation Mechanisms

- **Cross-Functional Teams:** Form interdisciplinary teams responsible for evaluating and overseeing GenAI procurement and deployment.
- **Pilot Programs:** Begin with small-scale pilots or sandbox environments to test GenAI solutions rigorously before scaling.
- **Structured Procurement Processes:** Utilize performance-based contracting methods, integrating AI-specific evaluation criteria into RFPs and contracts.
- **Continuous Training:** Implement specialized training programs to educate procurement and technical staff on AI risks, ethical implications, and procurement processes.
- **Security and Compliance Checks:** Ensure GenAI vendors meet established cybersecurity and data privacy requirements (e.g., FedRAMP compliance for cloud-based solutions).
- **Regular Monitoring:** Set up processes for continuous monitoring of AI performance, ethical implications, and cost management, with periodic reassessment and validation.
- **Flexible Procurement Models:** Utilizing agile, modular, phased, and challenge-based procurement methods, as well as flexible vendor pools and cooperative agreements, to enable rapid technological adaptation and significantly reduce procurement cycles from years to months or weeks.

Next Steps for Caltrans

- **Immediate Actions:** Develop a detailed GenAI procurement framework incorporating existing state guidance, federal recommendations, and successful state models like California's RFI², including agile and iterative procurement processes.
- **Stakeholder Engagement:** Convene internal stakeholders from IT, procurement, privacy, and legal departments to establish unified procurement protocols.
- **Pilot Identification:** Identify specific, actionable use cases within Caltrans operations suitable for initial GenAI pilot projects.
- **Governance Structure Setup:** Establish governance and oversight structures, including an AI Governance Board, to guide implementation and ethical compliance.
- **Capacity Building:** Launch training sessions for procurement officers and technical teams focusing on GenAI procurement best practices, risk management, and ethical considerations.
- **Performance Metrics Development:** Define clear metrics for evaluating GenAI performance, ensuring they align with Caltrans' operational goals and ethical standards.

Detailed Findings

Generative AI (GenAI) technologies offer transformative potential for government operations, from automating tedious tasks to uncovering new insights. Public agencies – including state Departments of Transportation (DOTs) – are exploring GenAI to improve services like traffic management, customer service, and data analysis. At the same time, officials recognize significant risks around bias, accuracy, security, and accountability. In late 2023, *Executive Order 14110*[1] *on Safe, Secure, and Trustworthy AI* catalyzed a flurry of activity to ensure responsible adoption of AI across government. Since then, federal and state agencies have been developing policies, frameworks, and practical guidance to navigate GenAI procurement in a way that captures its benefits while managing its risks[2]. While considerations around bias, accuracy, and security remain relevant, the federal policy landscape has shifted. The January 2025 Executive Order Removing Barriers to American Leadership in Artificial Intelligence[3] refocuses government efforts on accelerating innovation, reducing regulatory friction, and promoting U.S. leadership in AI.

This report examines the current and recommended best practices for GenAI procurement in the public sector, focusing on four dimensions: (1) procurement policy frameworks, (2) vendor selection criteria (including risk, fairness, and public-interest alignment), (3) data privacy and security requirements, and (4) implementation and integration strategies (piloting, adoption, and scaling). Throughout, we highlight case studies and guidance from government entities (e.g. U.S. General Services Administration, National Institute of Standards and Technology, Office of Management and Budget) and professional associations (e.g. NASCIO, AASHTO), with an emphasis on U.S. state and federal examples.

AI is also advancing faster than traditional government procurement can keep up. State DOTs and agencies are shifting to agile, faster models to deliver tech solutions in months, not years. New approaches like automation, AI, and challenge-based procurement (e.g., California's RFI²) aim to streamline processes, cut costs, and improve services. A separate section explores how agencies are modernizing procurement to keep pace with innovation.

Procurement Policy Frameworks for GenAI

Public-sector procurement of GenAI is being guided by emerging policies at both federal and state levels. The White House Office of Management and Budget (OMB) in October 2024 released a government-wide guidelines for acquiring AI software, emphasizing a culture of risk management and interagency collaboration[4]. Key provisions include early involvement of privacy officials in the acquisition process, cross-functional AI governance boards, and use of outcomes-based procurement techniques to keep pace with rapid technology change. At the most basic level, because generative AI tools are software, agencies casn treat GenAI acquisitions no differently than other IT procurements with respect to existing laws and regulations, while layering on additional AI-specific diligence[5]. In other words, GenAI tools are still "software" subject to standard IT acquisition rules (competition requirements, security reviews, etc.), but agencies must also address the unique challenges of AI (e.g. algorithmic bias, "hallucinations," and model transparency) within those procurement processes.

[1] https://bidenwhitehouse.archives.gov/omb/briefing-room/2023/11/01/omb-releases-implementation-guidance-following-president-bidens-executive-order-on-artificial-intelligence/

[2] https://gsablogs.gsa.gov/technology/2024/04/29/generative-ai-and-specialized-computing-infrastructure-acquisition-resource-guide-now-available/

[3] https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence/

[4] https://www.whitehouse.gov/wp-content/uploads/2024/10/M-24-18-AI-Acquisition-Memorandum.pdf

[5] https://gsablogs.gsa.gov/technology/2024/04/29/generative-ai-and-specialized-computing-infrastructure-acquisition-resource-guide-now-available/

The U.S. General Services Administration (GSA) has published a *Generative AI Acquisition Resource Guide* (April 2024). This GSA guide, produced in response to Executive Order 14110[6]'s mandate, lays out common use cases, risk considerations, and questions contracting officers should ask when scoping GenAI procurements[7]. It recommends defining the agency problem *before* picking a solution and suggests using testbed environments or sandboxes to pilot GenAI tools on a small scale before any large-scale acquisition.

A new version of OMB memorandum (M-25-22, "Driving Efficient Acquisition of AI in Government"[8]) requires agencies to integrate AI considerations into their procurement policies. The memo is guided by three core goals: promoting a competitive U.S. AI marketplace, protecting taxpayer dollars through performance tracking and risk management, and enhancing acquisition effectiveness via cross-functional collaboration. Agencies must revise acquisition procedures within 270 days and are encouraged to adopt performance-based contracting, prevent vendor lock-in, and implement continuous testing and monitoring of AI systems. The memo emphasizes responsible data use, directing agencies to protect government data rights while respecting contractor IP, and to restrict vendors from using government data to train commercial AI models without explicit consent. It also tasks the GSA with developing procurement guides and a resource-sharing platform[9]. Additionally, agencies are advised to consider requiring contractors to disclose any use of AI to ensure transparency and mitigate risks. Overall, federal policy is coalescing around a framework that blends existing acquisition rules (Federal Acquisition Regulation, IT acquisition policies) with AI risk management practices, guided in part by NIST's AI Risk Management Framework (RMF)[10] as a baseline for trustworthy AI.

State governments are likewise advancing GenAI procurement governance. For example, California's Governor issued an *Executive Order on AI* (Sept 2023)[11] directing state agencies to develop guidelines for public-sector procurement and use of generative AI based on NIST's AI RMF principles. In response, California's Government Operations Agency in collaboration with other state agencies released comprehensive GenAI guidelines for Public Sector Procurement[12], which draw on the NIST RMF's concepts of mapping AI risks, measuring and managing those risks, and governing AI use in a continuous cycle. California stood up a dedicated portal (genai.ca.gov[13]) to assist state teams with GenAI initiatives "from training to procurement". This includes step-by-step playbooks for different scenarios – e.g. discovering GenAI features in an existing vendor product, buying a new GenAI tool, or renewing a contract that adds GenAI capabilities[14].

Other states have taken similar actions: Washington State Governor's Executive Order 24-01[15] led to the creation of initial procurement guidelines for generative AI[16] (published by WaTech in 2024) that agencies must follow when evaluating, piloting, or deploying GenAI technologies. Washington's guidance explicitly states that it *"shall defer to and work in concert with"* all existing state procurement statutes and enterprise IT policies, reinforcing that AI procurements must not bypass foundational procurement rules. It then augments those rules with AI-specific process checkpoints (detailed further below in vendor selection and risk management). Meanwhile, states like Idaho have established AI task forces (the Idaho AI Advisory Council, created by a 2024 bill) to monitor and assess agency AI deployments with an emphasis on

[6] https://bidenwhitehouse.archives.gov/omb/briefing-room/2023/11/01/omb-releases-implementation-guidance-following-president-bidens-executive-order-on-artificial-intelligence/

[7] https://www.gsa.gov/about-us/newsroom/news-releases/gsa-releases-generative-ai-acquisition-resource-gu-04292024

[8] https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-22-Driving-Efficient-Acquisition-of-Artificial-Intelligence-in-Government.pdf

[9] https://www.gsa.gov/technology/government-it-initiatives/artificial-intelligence/procuring-ai-solutions

[10] https://www.fedramp.gov/2024-06-27-release-of-et-framework/

[11] https://www.gov.ca.gov/wp-content/uploads/2023/09/AI-EO-No.12-_-GGN-Signed.pdf

[12] https://www.govops.ca.gov/wp-content/uploads/sites/11/2024/03/3.a-GenAI-Guidelines.pdf

[13] https://www.genai.ca.gov/

[14] https://www.genai.ca.gov/choose-your-journey/

[15] https://governor.wa.gov/sites/default/files/exe_order/24-01%20-%20Artificial%20Intelligence%20%28tmp%29.pdf

[16] https://watech.wa.gov/sites/default/files/2024-11/Initial%20Procurement%20Guidelines%20for%20GenAI%20Final.pdf

transparency and accountability[17]. Across jurisdictions, there is a trend toward standing up AI governance structures (chief AI officers, interagency AI councils, etc.) and issuing "responsible AI" guidelines that encompass procurement. The National Association of State CIOs (NASCIO) has even released an *AI Blueprint* (December 2023) outlining 12 key considerations to help states develop AI roadmaps, with the goal that "2024 will be all about AI governance" in the public sector[18]. These considerations reinforce the need for clear procurement protocols, ethical use policies, and workforce training as states ramp up AI use.

Vendor Selection Criteria and Contracting Considerations

Selecting vendors and solutions for generative AI involves new evaluation criteria to address algorithmic risks, ethics, and performance, layered atop traditional factors like cost and technical fit. Best practices recommend a multi-disciplinary procurement team approach: involve not only procurement officers, but also IT, security, legal, and program experts *early* in drafting requirements and evaluating proposals[19]. For instance, OMB now requires agencies to include an internal cross-functional team (and by extension, data security officials) for any AI acquisition to help identify privacy risks and mitigation strategies[20]. Many agencies are also appointing Chief AI Officers or AI review boards to advise on procurements, ensuring that those with AI expertise assess vendor claims and risk controls[21]. By assembling a cross- functional team, governments can better evaluate a GenAI solution's fitness and trustworthiness before purchase.

Risk and Fairness Assessments

A core vendor selection criterion is the risk profile of the GenAI system. Leading guidance calls for determining whether the use-case is "high-impact" – e.g. if the AI's output will significantly affect individuals' rights, safety, or access to services[22]. If so, extra scrutiny is warranted. Agencies like NIST and OMB advise conducting formal AI risk assessments (sometimes called Algorithmic Impact Assessments) prior to award for high-impact AI systems. This includes evaluating: the intended purpose and benefits of the tool, the quality of its training data, potential biases or harm its outputs could produce, plans for human oversight, and how the vendor will enable ongoing performance monitoring[22].

In practice, procurement teams should require vendors to demonstrate their AI's reliability and fairness. For example, OMB's procurement memo instructs agencies to *test* any offered AI solutions during the evaluation process to understand their capabilities and limitations[23]. Failure to report GenAI may result in disqualification[24]. If a bidder provides a GenAI model, the agency might run sample scenarios or trial runs (in a sandbox) to verify claims about accuracy, check for "hallucinated" outputs, and observe how the model handles edge cases.

Bias mitigation is a critical factor: agencies should ask vendors what steps they have taken to prevent discriminatory outcomes (e.g. biases against protected classes) and whether they have tested the model for fairness. Washington State's GenAI guidance, for instance, includes specific vendor questions like *"What quality control is in place to test and monitor for potential bias?"* and *"Describe any public or community engagement used in the system's design – has input from affected communities been*

[17] https://statescoop.com/generative-ai-government-procurement-risk-2024/
[18] https://www.govtech.com/artificial-intelligence/nascio-resource-offers-states-12-key-considerations-for-ai
[19] https://arxiv.org/pdf/2411.04994v1
[20] https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-22-Driving-Efficient-Acquisition-of-Artificial-Intelligence-in-Government.pdf
[21] https://www.workforcebulletin.com/new-federal-agency-policies-and-protocols-for-artificial-intelligence-utilization-and-procurement-can-provide-useful-guidance-for-private-entities
[22] https://www.workforcebulletin.com/new-federal-agency-policies-and-protocols-for-artificial-intelligence-utilization-and-procurement-can-provide-useful-guidance-for-private-entities
[23] https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-22-Driving-Efficient-Acquisition-of-Artificial-Intelligence-in-Government.pdf
[24] https://www.dgs.ca.gov/PD/Resources/SCM/TOC/23/2302

considered?"[25]. These questions aim to surface how the AI was developed and whether its use has been vetted for societal impact. Alignment with public values is increasingly important – vendors may be asked to affirm that their AI systems will not violate civil rights or privacy laws and that they align with ethical AI principles (such as the White House's AI Bill of Rights[26]).

Another emerging consideration is fairness in the procurement process itself when using AI. Agencies experimenting with AI-assisted vendor selection must be cautious that these tools don't introduce bias against certain vendors. For example, if a generative AI tool is used to scan and rank vendor proposals or suggest vendors, it might inadvertently favor firms that resemble past successful contractors (potentially disadvantaging minority-owned or small businesses). To counter this, procurement professionals need training to spot and mitigate biases in AI-generated suggestions[27]. In practice, this means any AI tool used in the RFP or bid evaluation process should be used *advisably* – as a helper, not an automated decider – and its outputs should be reviewed for fairness and accuracy by humans.

Evaluation Criteria and Contract Terms

In evaluating GenAI vendors, agencies are expanding RFP criteria and contract requirements to address:

- **Transparency and Explainability:** Does the vendor provide adequate documentation about how the AI system works? OMB urges agencies to obtain documentation that *"facilitates transparency and explainability"* of AI offerings[28]. Vendors should ideally disclose the provenance of training data, model architectures, and known limitations. Some solicitations include requirements for access to the model (or its outputs) for independent audits or explainability assessments, especially if decisions made by the AI impact citizens. Vendors who refuse basic transparency (e.g. treating the model as a complete black box) may be scored lower on proposals, particularly in high-stakes use cases.

- **Data Source and Quality:** Agencies want to know *what data the GenAI was trained on*, and whether it's appropriate for the intended use. For example, a state DOT considering a GenAI traffic management tool will ask if the model was trained on relevant transportation datasets or merely general internet data. Washington's sample questionnaire explicitly asks for the model name/version and *"what data is the model trained on and what data will be input into the system when in use?"*[29]. This informs evaluators about potential biases (training data gaps) and compliance (e.g. if any training data is proprietary or sensitive). Preference may begiven to vendors who can use agency- provided data securely to fine-tune models, thereby improving relevancy while protecting data control.

- **Performance and Accuracy:** Traditional performance metrics (throughput, latency) are now coupled with AI-specific metrics. Agencies may require vendors to report measures like the model's accuracy on relevant tasks, its false positive/negative rates, or other quality benchmarks. Washington's guidance asks *"What performance metrics were selected to determine the model's effectiveness? What is the model optimizing for, and under what constraints?"*[28]. During evaluation, agencies might run a standardized set of queries or test cases to compare vendors' GenAI outputs. Selecting a vendor often hinges on finding an AI whose output quality meets the agency's needs with minimal unacceptable errors.

- **Security and Compliance:** Vendor selection criteria include whether the AI solution and the vendor meet government security standards (addressed in the next section on data security). But from an

[25] https://watech.wa.gov/sites/default/files/2024-11/Initial%20Procurement%20Guidelines%20for%20GenAI%20Final.pdf
[26] https://bidenwhitehouse.archives.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf
[27] https://statescoop.com/generative-ai-government-procurement-risk-2024/
[28] https://www.workforcebulletin.com/new-federal-agency-policies-and-protocols-for-artificial-intelligence-utilization-and-procurement-can-provide-useful-guidance-for-private-entities
[29] https://watech.wa.gov/sites/default/files/2024-11/Initial%20Procurement%20Guidelines%20for%20GenAI%20Final.pdf

evaluation standpoint, agencies will check if the vendor has necessary security certifications (e.g. FedRAMP authorization for cloud services) and compliance with relevant regulations (like HIPAA[30] if health data is involved, or state privacy laws). OMB's guidance specifically introduces new contract negotiation requirements to safeguard government interests: contracts must address privacy and security, and agencies should assess how vendors handle government data and intellectual property during proposal evaluations[27].

- **Avoiding Vendor Lock-In:** A notable procurement best practice is to avoid being trapped with a single AI provider. OMB's memo encourages structuring contracts and evaluation criteria to promote a competitive AI marketplace and prevent lock-in[27]. This might involve favoring solutions built on open standards or requiring that models be portable or that the agency retains rights to its data and any derived model improvements. In fact, M-25-22[31] instructs agencies to scope IP and licensing rights such that critical components (e.g. model weights, training data, or at least the outputs and logs) remain accessible to the agency even if the contract ends. During vendor selection, proposals are reviewed for how they address exit strategies – e.g. can the system be transitioned to another platform or internal system if needed? Vendors that offer robust exit clauses, data transfer assistance, or escrow of AI models[32] may score better on this dimension.

Finally, effective vendor selection for GenAI often involves a trade-off analysis between innovation and risk. Public agencies are encouraged to seek out innovative solutions (including from non-traditional gov contractors or startups), to promote a competitive AI market with innovative acquisition, but without compromising on ethical and security standards. This requires well-defined evaluation frameworks. Some organizations may use scorecards or rubrics explicitly incorporating the above criteria (transparency, bias mitigation, data usage, etc.), sometimes weighting them equally to price and technical capability. By embedding these criteria into RFPs and contracts, governments align vendor incentives with public sector values – pushing AI vendors to build more accountable and fair AI systems if they want government business.

Data Privacy and Security Requirements

Protecting sensitive data and ensuring the security of AI systems are paramount when government agencies procure GenAI technologies. Generative AI solutions typically rely on large volumes of data (both for training and as input during use), which raises unique privacy and security considerations. Best practices in the public sector establish clear safeguards and requirements in this area:

- **Privacy Compliance and Oversight:** Agencies must ensure that GenAI acquisitions comply with all privacy laws and regulations (e.g. the Privacy Act, GDPR for any personal data, state privacy statutes). OMB's guidance explicitly requires agencies to establish policies and processes to ensure compliance with privacy requirements in law and policy for AI systems[33]. In practical terms, this means any procurement involving personal or sensitive data should trigger a Privacy Impact Assessment and review by the agency's privacy officials. As noted, OMB now mandates privacy officer involvement throughout the AI procurement to identify risks of exposing PII (Personally Identifiable Information) or violating privacy rights[31].

  A common safeguard is to minimize data collection: contract terms should stipulate that the AI vendor will only collect or use government data when *reasonably necessary* for the contract's

[30] https://www.hhs.gov/hipaa/for-professionals/privacy/laws-regulations/index.html
[31] https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-22-Driving-Efficient-Acquisition-of-Artificial-Intelligence-in-Government.pdf
[32] https://blog.codekeeper.co/protecting-ai-with-software-escrow
[33] https://www.workforcebulletin.com/new-federal-agency-policies-and-protocols-for-artificial-intelligence-utilization-and-procurement-can-provide-useful-guidance-for-private-entities

purposes[31]. Any data not needed should remain with the agency. For example, if a DOT uses a GenAI tool for analyzing roadway images, the contract might forbid the vendor from copying the entire image database, instead allowing access via a secure API. Privacy protections also extend to output data – if the AI generates content based on sensitive inputs, those outputs may themselves be sensitive (e.g. an AI-generated summary of a confidential report) and must be handled accordingly. Agency chief data officers and counsel often insert clauses ensuring that data outputs are treated with the same confidentiality as inputs, and that individuals' privacy rights are respected in AI-driven decisions (aligning with principles from the AI Bill of Rights[34] like data privacy and notice).

- **Restrictions on Data Use and Model Training:** A critical requirement emerging in GenAI procurements is that vendors cannot use government data to train their models for other purposes without permission. Generative AI providers often improve their models by learning from user inputs and outputs – but when the "user" is a government agency, this poses a risk of sensitive data leakage. To address this, OMB's procurement memo directs agencies to include contract terms that permanently prohibit the use of non-public inputted and outputted results to further train a vendor's publicly or commercially available AI algorithms absent explicit agency consent[35]. In short, unless the government explicitly agrees, the vendor may not feed the agency's data or queries back into its broader model. Many early contracts for GenAI (such as those for using large language model APIs) now contain no-feed clauses or data isolation provisions.

  State agencies are following suit: California's guidelines[36] instruct procurement officers to ensure similar language[37]. This protects citizen data and government-sensitive information from inadvertent exposure or intellectual property loss. Agencies also demand documentation of data handling – vendors should describe how data will flow, be stored, encrypted, and eventually deleted or returned. Some contracts require a data management plan from the vendor as part of the proposal, which is evaluated by security and privacy teams[38].

- **Cybersecurity Standards:** GenAI tools must meet the same security standards as any IT system handling government data. As GSA emphasizes, contracting officers should consider cybersecurity, supply chain risk management, data governance and other standards and guidelines just as they would with other IT procurements[39]. This means if a GenAI solution is cloud-based (software-as-a-service), it likely needs to be FedRAMP authorized at the appropriate impact level (Moderate or High, depending on the data).

  Recognizing the surge in AI services, the FedRAMP program in 2023-2024 created an "Emerging Tech" fast-track process to expedite approvals for generative AI cloud offerings[40]. Agencies are advised to use FedRAMP marketplace solutions or insist that vendors obtain FedRAMP authorization. On-premise AI systems similarly must align with NIST security controls (800-53)[41] and agency-specific security policies.

  Procurement solicitations often include a Security Requirements Appendix covering topics like access control, encryption (for data at-rest/in- transit), incident response, and vulnerability testing. GenAI introduces a few special security concerns – for example, model integrity (ensuring the AI model

[34] https://bidenwhitehouse.archives.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf
[35] https://www.workforcebulletin.com/new-federal-agency-policies-and-protocols-for-artificial-intelligence-utilization-and-procurement-can-provide-useful-guidance-for-private-entities
[36] https://www.genai.ca.gov/choose-your-journey/incidental/follow-procurement-process/procurement-process/
[37] https://www.genai.ca.gov/choose-your-journey/incidental/follow-procurement-process/disclosure-contract-language/
[38] https://www.genai.ca.gov/choose-your-journey/incidental/monitoring-product-roadmap/contract-monitoring-evaluation/
[39] https://gsablogs.gsa.gov/technology/2024/04/29/generative-ai-and-specialized-computing-infrastructure-acquisition-resource-guide-now-available/
[40] https://www.gsa.gov/blog/2024/06/27/making-access-to-ai-a-priority
[41] https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r5.pdf

hasn't been tampered with or poisoned with malicious data)[42]. While still an emerging area, agencies might require the vendor to provide hashes or provenance for the model version, or to undergo independent evaluation for adversarial robustness[43] if the use is mission-critical (e.g. AI used in transportation safety systems should be tested against adversarial inputs). Supply chain risk management (SCRM)[44] is also emphasized: agencies will check where the AI model was developed and hosted – foreign adversary concerns mean some will avoid vendors whose models or infrastructure are tied to high-risk countries or entities, in line with supply chain security regulations.

- **Data Protection and Secure Environments:** Government buyers are encouraged to insist on secure hosting and data segregation for GenAI. For instance, if a state DOT contracts a vendor to provide a GenAI analytics platform, the contract may require that the solution be hosted in a Government Community Cloud or a tenant isolated to that agency, with rigorous identity management. Some states have stood up central AI sandboxes or secure environments where agencies can experiment with GenAI on sanitized data[45]. The U.S. DOT's Transforming Transportation Advisory Committee (TTAC)[46] in late 2024 recommended providing a set of enterprise-approved generative AI tools (such as a version of Microsoft's CoPilot) to all staff in a controlled manner. This allows employees to use GenAI with certain guardrails (logging, filtering of sensitive content, etc.), rather than turning to unsanctioned public tools. On the procurement side, this means if an agency is buying a GenAI tool, it may prefer vendors that can deploy in the agency's own cloud environment or that have robust admin controls. Additionally, continuous monitoring requirements are common: contracts might stipulate that the vendor will support security assessments, provide audit logs of AI usage, and promptly report any data breaches or model misbehavior. OMB's AI use guidance calls for ongoing monitoring of AI performance and impacts[47], which intersects with security: for example, unusual spikes in AI output errors might indicate a data drift or security issue. Agencies should plan for periodic re-validation of the AI system's accuracy and security post-deployment, and vendors are typically obligated to patch or improve the model to address any identified vulnerabilities or unacceptable outcomes.

- **Legal and Intellectual Property Safeguards:** Data-related contract clauses also cover intellectual property (IP) rights and liability. Agencies are ensuring that they retain appropriate rights to any data outputs or AI-generated content. For example, if a GenAI tool produces a new engineering report or code, the contract might clarify who owns that output – often, the government will assert rights (at least a license) to use and modify the AI-generated work product. OMB's memo suggests standardizing such clauses and considering IP rights based on the intended use of AI, to avoid vendor lock-in[45]. Liability and indemnity clauses are being updated for AI: vendors may be asked to indemnify the agency if the AI outputs infringe on someone's copyright or if a data breach occurs on the vendor's side. From a privacy standpoint, contracts also address remedy for individual harms – for instance, if an AI system mistakenly exposes personal data or makes an adverse decision about a citizen, what is the recourse? While procurement contracts can't solve all governance questions, forward-looking agencies are tying contract performance to compliance with AI ethics guidelines[48] [49] [50]. A vendor might be contractually required to abide by the agency's AI ethics policy

[42] https://al-kindipublishers.org/index.php/jcsts/article/view/9461/8114
[43] https://datascientest.com/en/all-about-adversarial-robustness
[44] https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-161r1.pdf
[45] https://innovate-us.org/blog/state-and-local-governments-share-their-journey-to-build-ai-sandboxes
[46] https://www.transportation.gov/sites/dot.gov/files/2025-01/TTAC%202024%20Report.pdf
[47] https://www.workforcebulletin.com/new-federal-agency-policies-and-protocols-for-artificial-intelligence-utilization-and-procurement-can-provide-useful-guidance-for-private-entities
[48] https://www.intelligence.gov/ai/ai-ethics-framework
[49] https://www.ai.mil/Initiatives/Responsible-AI/
[50] https://department.va.gov/ai/trustworthy-ai/

or risk termination.

In summary, public agencies are building robust privacy and security protections into GenAI procurements. The mantra is "protect the data, protect the people": ensure the AI system doesn't expose sensitive information, and that it operates within a secure, legally compliant boundary. As a case in point, the U.S. DOT's TTAC[51] explicitly advised USDOT to build responsible use awareness among users of generative AI tools, especially regarding safe and secure data management practices. Responsible AI procurement thus goes hand-in-hand with responsible data stewardship. By setting these clear requirements upfront and verifying vendor practices, government agencies can harness GenAI innovations while upholding public trust and security.

Implementation and Integration Strategies

Procuring a generative AI tool is only the first step – agencies must also pilot, implement, and scale these technologies in a way that maximizes value and minimizes disruption. Best practices stress a careful, phased approach to GenAI adoption within government operations:

- **Startwith Clear Needs and Small Pilots:** Rather than deploying AI broadly right away, agencies are urged to identify specific business needs or problems that GenAI can address, and begin with limited-scope pilots. The GSA guide emphasizes beginning *with* your agency's needs. Rather than starting with solutions and specifications, define the problem that the agency wants generative AI tools to help solve[52]. This ensures technology is driven by mission goals (e.g. automating report drafting, improving citizen inquiry response, analyzing traffic data) and avoids the trap of adopting AI for hype's sake. Once a use-case is defined, agencies should conduct proof-of-concepts or sandbox tests. Because GenAI tools are evolving rapidly, GSA recommends using testbeds and sandboxes to try solutions before committing to large scale buys with too many unknowns[50]. For example, a DOT might pilot a GenAI model on a subset of transportation data to see if it can accurately predict accident hotspots.

  Pilot programs help evaluate performance, uncover unexpected behaviors, and refine requirements before a full procurement. Indeed, some state frameworks build this step in: Washington's guidance[53] includes a "Procurement/Development" phase where iterative testing occurs, and California's playbook explicitly has a stage to "perform market research, inventory your data, and assess your data readiness" followed by a stage to do risk assessment and prepare procurement package– effectively ensuring a trial and learning period precedes any major acquisition[54].

- **Form Cross-Functional Teams and Governance:** Implementing GenAI is not solely an IT project; it requires interdisciplinary governance. Agencies are advised to form a project team that includes the program owners (business unit), IT experts, data scientists (if available), procurement officials, legal/policy advisors, and ethics or equity advisors. California's GenAI adoption framework has an early step to "form your team and develop department challenge"[55], ensuring all stakeholders coalesce around the project's objectives and constraints. This team should develop clear success metrics for the pilot (e.g. reduce processing time by X%, maintain error rate below Y, etc.) and decide how human oversight will be integrated.

  Many agencies also require any AI project to go through an AI governance board or get approval

[51] https://www.transportation.gov/sites/dot.gov/files/2025-01/TTAC%202024%20Report.pdf
[52] https://www.gsa.gov/technology/government-it-initiatives/artificial-intelligence/procuring-ai-solutions
[53] https://watech.wa.gov/sites/default/files/2024-11/Initial%20Procurement%20Guidelines%20for%20GenAI%20Final.pdf
[54] https://www.genai.ca.gov/choose-your-journey/procure-tool/take-training/
[55] https://www.genai.ca.gov/choose-your-journey/procure-tool/form-team-develop-challenge/

from the CIO/Chief AI Officer at certain checkpoints. For instance, California mandates that if a GenAI project is deemed moderate- or high- risk (using criteria akin to NIST's impact levels), the department must consult with the California Department of Technology (CDT) for guidance or approval[56]. This central oversight helps maintain standards and share lessons learned across agencies. Similarly, OMB's guidance pushes agencies to convene AI Governance Boards[57].

- **Incorporate Ethical and Equity Considerations:** A notable best practice is to address equity and ethics from the design phase. Generative AI can unintentionally perpetuate biases or produce outcomes that negatively impact vulnerable communities if not checked. Some agencies have published *Guidelines for Evaluating Impacts of Generative AI on Vulnerable and Marginalized Communities*[58], and their procurement playbook makes "Consider equity impacts" a distinct step before issuing an RFP[59]. Practically, this means asking questions like: Who could be harmed or excluded by this AI system? Does the model perform equally well for different demographic groups? Are there historical biases in the data that need mitigation? Some pilot projects include an ethics review or consult with community stakeholders to get these answers. For example, if a state deploys a GenAI chatbot for public services, they might involve racial equity offices or disability advocates to ensure the tool's responses are inclusive and accessible. This early consideration can inform requirements (e.g. demanding the vendor to demonstrate bias testing results) and set up appropriate human-in-the-loop controls during implementation (so humans can override or correct AI outputs that seem problematic). The aim is to align GenAI implementations with public sector values of fairness and non-discrimination from the outset.

- **Leverage Training and Change Management:** Integrating GenAI into government workflows requires preparing the workforce. Training programs for staff – both the technical implementers and the end-users – are highly recommended. California, for instance, has developed GenAI training modules for state employees and a specialized procurement training through its Procurement and Contracting Academy (CalPCA) focused on GenAI[60]. This training covers how to identify GenAI in vendor offerings, understanding risk levels, and following the state's GenAI procurement process. On a broader scale, agencies are encouraging a culture of learning around AI. The U.S. DOT's advisory committee suggested that staff should "embrace curiosity, candor, competence, and caution when using AI tools," and that organizations provide safe environments for employees to experiment and build competence in AI[61].

  Change management efforts – such as workshops, pilot user feedback sessions, and clear communication from leadership – help ease employee concerns ("Will AI take my job?" or "Am I allowed to use ChatGPT?") and foster acceptance. Some agencies have issued internal guidelines on appropriate use of GenAI (for example, Vermont and others have "State Employee Use of Generative AI" policies) to both permit innovation and set boundaries (like not uploading sensitive files to public AI tools)[62] [63]. Successful implementation often hinges on employees understanding both the capabilities and the limits of GenAI, so they can use it as an augmenting tool rather than a magic solution.

- **Phased Deployment and Scaling:** After a successful pilot, scaling GenAI should be done iteratively. Procurement strategies can facilitate this by starting with a pilot contract or phase 1, with options to

[56] https://www.genai.ca.gov/choose-your-journey/procure-tool/consult-cdt/
[57] https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-21-Accelerating-Federal-Use-of-AI-through-Innovation-Governance-and-Public-Trust.pdf
[58] https://www.genai.ca.gov/wp-content/uploads/sites/360/2024/12/GenAI-Equity-Guidelines-2024.12.19.pdf
[59] https://www.genai.ca.gov/choose-your-journey/procure-tool/consider-equity-impacts/
[60] https://www.genai.ca.gov/choose-your-journey/procure-tool/take-training/
[61] https://www.transportation.gov/sites/dot.gov/files/2025-01/TTAC%202024%20Report.pdf
[62] https://digitalservices.vermont.gov/sites/digitalservices/files/documents/Guidelines%20for%20use%20of%20Generative%20AI.pdf
[63] https://www.dgs.ca.gov/Resources/SAM/TOC/4900/4986-12

expand. Many agencies utilize incremental contracting – e.g. a contract that has a base period for a small implementation, and additional phases that are exercised only if performance is satisfactory and funds are approved for expansion (e.g., modernize management of the U.S. Army's Test and Evaluation business processes[64]). This allows the organization to evaluate real- world results and impacts before full commitment. As the solution scales, agencies should update policies and controls accordingly. For instance, a state DOT might pilot an AI model in one region or for one type of analysis, then, if results are good, gradually roll it out statewide with necessary adaptations (like more computing resources or additional training data for new contexts). During scaling, continuous monitoring is crucial: performance metrics and error rates should be tracked to ensure the AI continues to operate as expected with larger scope. OMB's guidance underlines that agencies must conduct ongoing monitoring for performance and potential adverse impacts throughout the AI system's lifecycle[65]. A state DOT might plan regular recertification of AI tools – for example, an annual review of a GenAI application's outputs, its security posture, and whether any drift or new risks have emerged. This ties back into procurement via contract administration: contracts for AI solutions often include requirements for periodic performance reports, security audits, and the flexibility to make updates. OMB even suggests that at contract closeout, agencies work with the vendor to ensure continued access to data or models if needed (avoiding a scenario where the agency loses functionality at contract end)[66].

- **Cost Management:** Implementing GenAI can incur unpredictable costs, especially with cloud-based models that charge per use or by volume of data processed. A best practice is to monitor usage closely and put guardrails on spending. GSA cautions that generative AI is very often billed like other Software as a Service so usage costs can really grow quickly if not appropriately monitored and managed[67]. During pilots, agencies should gather data on how much a given use-case consumes (e.g. API calls, tokens, compute hours) and project the budget for scale. Procurement can address this by negotiating pricing models (enterprise licenses vs. pay-as-you-go, bulk discounts after a certain number of queries, etc.) and by implementing controls like quotas or budget caps in the software. Some also compare the cost-benefit: e.g. if a GenAI tool saves staff 1000 hours of work but costs $X in cloud fees, is it justified? Having this analysis helps in making informed decisions about scaling up or seeking alternative solutions. In practice, agencies often start with a limited budget pilot and only scale when they are confident in the Return on Investment (ROI) and have funding allocated.
- **Integration with Existing Systems:** Another implementation aspect is technical integration. GenAI tools should be integrated into agencies' existing IT infrastructure and workflows in a secure and user-friendly way[68]. This may involve building APIs between the GenAI service and legacy databases or embedding an AI assistant into existing software (like integrating a GPT-based chatbot into the DOT's customer service portal). Agencies should involve their IT architecture teams to ensure compatibility and to avoid creating silos. Many vendors provide integration support, but agencies need to plan for it – sometimes issuing a task order to an IT integrator to handle customizing and integrating the AI solution. It's also wise to maintain a human fallback: for example, if a generative AI fails to answer a citizen's question in a chatbot scenario, the system should seamlessly route the query to a human agent. Implementation plans should map out such

[64] https://bigbear.ai/newsroom/bigbear-ai-selected-for-9-month-7-7m-ota-contract-as-the-single-provider-for-phase-2-u-s-army-test-and-evaluation-command-integrated-mission-management-system/
[65] https://www.workforcebulletin.com/new-federal-agency-policies-and-protocols-for-artificial-intelligence-utilization-and-procurement-can-provide-useful-guidance-for-private-entities
[66] https://www.dgs.ca.gov/Resources/SAM/TOC/4900/4986-12
[67] https://gsablogs.gsa.gov/technology/2024/04/29/generative-ai-and-specialized-computing-infrastructure-acquisition-resource-guide-now-available/
[68] https://www.dhs.gov/sites/default/files/2025-01/25_0106_ocio_dhs-playbook-for-public-sector-generative-artificial-intelligence-deployment-508-signed.pdf

contingencies.

- **Ongoing Learning and Knowledge Sharing:** Governments are increasingly collaborating and sharing lessons on GenAI implementations. Communities of practice (like the AI Community of Practice[69], or NASCIO's AI working group[70]) facilitate exchange of experiences – what use cases worked, what pitfalls to avoid. For state DOTs, organizations such as AASHTO and the Transportation Research Board are beginning to host panels and produce reports on AI in transportation and implementation challenges[71]. AASHTO's knowledge sessions have noted that AI's role in transportation is growing and DOTs must be proactive in governance and skills development to capitalize on it. Following recommended strategies, some DOTs have launched innovative GenAI pilot projects: for example, California DOT (Caltrans) is piloting a GenAI tool to identify "potentially dangerous intersections" for pedestrians and cyclists, using AI to analyze patterns of incidents. Early results from such pilots can inform wider adoption (Caltrans can share if the GenAI correctly flagged high-risk sites and how it's being integrated into their safety interventions). Another pilot in California is exploring GenAI for improving call center responses in the Department of Tax and Fee Administration, aiming to provide accurate answers faster by having AI assist human agents[72]. These case studies illustrate a pattern: targeted GenAI projects with clear metrics (safety improvements, faster service) that, if successful, can be scaled up through formal procurement and replicated by other agencies.

In one striking local example of GenAI integration, the City of Murray, UT used ChatGPT to drastically speed up drafting a complex RFP (for a large holiday tree installation). What normally could take months was completed in a matter of days, with the proposal approved next-day and a vendor selected in 10 days[73]. The procurement officer credited generative AI with making it possible to meet the tight deadline, noting that it helped her translate stakeholder inputs into proper procurement language quickly. This success led the city to integrate GenAI into updating standard procurement documents (making them clearer about project expectations and budgets). However, there is a need for caution and training specially training procurement professionals to watch for AI- generated errors or biases[74]. This encapsulates the implementation balancing act: GenAI can yield efficiency gains in government processes, but agencies must implement it with due diligence, educating staff and maintaining oversight to ensure outcomes remain accurate, fair, and aligned with public objectives.

Key Steps for GenAI Adoption

In summary, a structured approach to implement and scale generative AI in government might involve:

- **Plan and Educate** – Build internal knowledge (train teams on GenAI fundamentals), define the problem and success criteria, and ensure leadership buy-in and ethical guidelines are in place.
- **Pilot and Evaluate** – Conduct small-scale pilots or proof-of-concepts in a sandbox. Use cross- functional teams to evaluate results against metrics (accuracy, bias, efficiency). Adjust the solution or approach based on findings.
- **Risk Assessment and Approval** – Perform thorough risk assessments (security, privacy, impact) and have oversight bodies review high-risk AI projects. Obtain necessary approvals before full deployment.
- **Procure for Scale** – If pilot is successful, proceed with a scalable procurement (e.g. a contract with

[69] https://coe.gsa.gov/communities/ai.html
[70] https://www.nascio.org/about/committeesandforums/#toggle-id-6
[71] https://aashtojournal.transportation.org/knowledge-session-examining-role-of-ai-in-transportation/
[72] https://www.genai.ca.gov/ca-action/projects/call-center-productivity/
[73] https://statescoop.com/generative-ai-government-procurement-risk-2024/
[74] https://www.nascio.org/about/committeesandforums/#toggle-id-6

expansion options). In RFPs/contracts, incorporate lessons from the pilot and include requirements for ongoing risk mitigation (transparency, data handling, etc.).

- **Integrate and Train Users** – Integrate the GenAI tool into existing systems and workflows. Provide training or change management for end-users so they know how to effectively use the AI tool and understand its limitations.
- **Monitor and Iterate** – Continuously monitor the AI system's performance, costs, and impacts. Establish feedback loops: collect user feedback, audit outputs periodically (for errors or bias), and require the vendor to update or tune the model as needed. If issues arise, be ready to pull back or implement additional controls. Document and share outcomes and best practices with other agencies (building a knowledge base for future AI procurements).

By following such strategies, government agencies can incrementally build confidence in generative AI technologies. Starting small and focused, involving the right experts, and scaling deliberately allows the public sector to innovate with AI carefully – gaining improvements in efficiency and service delivery while upholding the standards of fairness, transparency, and security that citizens expect. The ongoing efforts by GSA, OMB, NIST, NASCIO, and others to provide frameworks and playbooks will continue to mature.

Reimagining DOT Procurement for Rapidly Evolving AI Technology

Generative AI and related technologies are advancing at breakneck speed, while traditional government procurement cycles do not keep pace with fast-changing technology and user needs[75] [76]. State and local agencies including state Departments of Transportation (DOTs) are recognizing that this gap is unsustainable. As Idaho's Chief Information Officer put it, "Technology moves way too fast for us to be doing an 18 to 24 months procurement"[77]. To better keep pace with innovation, procurement processes themselves must be transformed. DOTs and other agencies are beginning to embrace nimble, agile procurement models that can deliver solutions in months instead of years. Procurement innovation and digital transformation also aim to streamline complex purchasing processes, reduce costs, and improve service delivery. Trends such as automation, AI, agile methods, sustainable sourcing, and challenge-based procurement (like the California's Request for Innovative Ideas (RFI²) model) are helping procurement professionals work more efficiently and adaptively[78]. This section explores emerging practices and recommendations – focusing on U.S. state DOTs and government agencies – to modernize procurement for both large-scale systems and small pilots or tools in the age of AI.

Emerging Procurement Innovations in Government

- **Agile and Iterative Procurement:** **Agile procurement**[79] applies principles from agile software development to contracting. Agile procurement means breaking down large, rigid contracts into smaller, flexible phases that allow agencies to work closely with vendors in iterative cycles. Instead of a long, fixed bidding and development process, the agency awards pilot contracts or task orders to quickly develop and test parts of a solution, gather feedback, and adjust requirements as needed. The key principles of agile procurement include delivering early value through incremental purchases, adapting quickly to change with flexible contracts, working in short cycles for faster feedback, focusing on lean sourcing to eliminate waste, and practicing continuous improvement via regular stakeholder input. These principles are supported by enabling technologies such as cloud-based platforms for real-time collaboration, data analytics

[75] https://www.usds.gov/report-to-congress/2016/procurement/
[76] https://www.cio.com/article/406456/government-tech-procurement-takes-three-times-longer-than-average.html
[77] https://statescoop.com/generative-ai-government-procurement-risk-2024/
[78] https://insider.govtech.com/california/sponsored/the-two-procurement-concepts-that-will-enable-california-to-make-the-most-of-taxpayer-dollars
[79] https://veridion.com/blog-posts/agile-procurement/

for performance and market insights, and AI tools that automate tasks like supplier sourcing and contract management (together making procurement faster, more adaptive, and more efficient).

The GAO Agile Assessment Guide[80] outlines several contracting strategies and best practices for implementing Agile procurement in the federal space. One key strategy is the use of modular contracting, which involves breaking large projects into smaller, independently managed components that can be acquired incrementally. This allows for greater flexibility and better alignment with Agile's iterative delivery model. Additionally, agencies are encouraged to use Statements of Objectives (SOOs) rather than traditional Statements of Work (SOWs) to allow vendors the freedom to propose creative, outcomes-based solutions. Flexible pricing models such as time-and-materials (T&M) contracts may also be more compatible with Agile than firm-fixed-price structures. Legal and institutional barriers are also acknowledged, with the Federal Acquisition Regulation (FAR) often viewed as a constraint. However, the guide clarifies that many Agile-friendly acquisition strategies are already permissible under existing FAR rules (the real challenge lies in limited awareness and risk aversion among contracting officers). A significant barrier is the lack of training for procurement professionals in Agile-specific practices, which impedes effective implementation. To overcome these challenges, the guide recommends several best practices. These include training both procurement and program staff in Agile principles, fostering collaboration across program, legal, and acquisition teams, and encouraging the use of challenge-based acquisition where vendors demonstrate capabilities before full-scale procurement. Agencies are also urged to adopt templates, case studies, and playbooks from resources such as 18F and the U.S. Digital Service. Finally, incremental funding strategies that align with Agile development cycles are advised to maintain flexibility and responsiveness throughout the project lifecycle.

The U.S. Department of Homeland Security's Procurement Innovation Lab and New Zealand's Dunedin City Council showcase effective use of agile procurement to streamline traditionally slow government processes. The DHS lab helped the FDA reduce contract award time from 180 to 94 days through a two-phase vendor selection process, involving short written proposals and oral presentations, which led to faster decisions and zero protests. Meanwhile, Dunedin City Council cut a three-month procurement process to just three days by shortlisting vendors quickly and holding a big room event for real-time solution presentations, enabling a final selection by day two. Both examples demonstrate how agile methods can accelerate procurement while maintaining transparency and vendor engagement[81].

- **Challenge-Based (Problem-First) Solicitations:** An increasingly popular model is challenge-based procurement, also called *problem-based* or *outcome- based* procurement. Instead of the traditional route of writing extremely detailed specifications for a predetermined solution, the agency defines the problem or challenge and invites the market to propose innovative solutions[82] [83]. This approach spurs creativity and flips the script (vendors become partners in designing the solution rather than just fulfillers of narrow requirements). California has been a leader in this space with its Request for Innovative Ideas (RFI²) process, established by Governor Newsom in 2019[84]. Under an RFI², the state publishes a problem statement instead of a strict RFP. Innovators then submit concept proposals, and the most promising ideas are selected for a funded proof-of-concept pilot – all within a single competitive procurement process (e.g., REQUEST FOR INNOVATIVE IDEAS

[80] https://www.gao.gov/assets/d24105506.pdf
[81] https://www.deloitte.com/content/dam/assets-zone1/tw/tc/docs/about/2025/government-trends-2024-en.pdf
[82] https://bloombergcities.jhu.edu/news/problem-based-procurement-crowding-new-solutions
[83] https://www.route-fifty.com/finance/2024/05/challenge-based-procurement-just-makes-sense/396942/
[84] https://www.gov.ca.gov/wp-content/uploads/2019/01/1.8.19-EO-N-04-19.pdf

(RFI²) #29442[85]). After demonstrating a working solution in the pilot, the top solution(s) can be scaled up through an implementation contract. In essence, it's a two-phase agile procurement: Phase 1 – solicit ideas and proof-of-concepts; Phase 2 – award the full solution contract to the best performer (e.g., REQUEST FOR INNOVATIVE IDEAS (RFI²) #29442). This challenge-based RFI² approach was used by Caltrans and California's Department of Technology in 2023–24 to solicit GenAI solutions for traffic safety and mobility, allowing the state to *test GenAI's potential* on a transportation problem.

Other governments are following suit. New York City, for example, is adopting challenge-based procurement modeled also[86]. Historically the city issued 1,000-page RFPs as if they know exactly all the contours of the problem, whereas now the city can issue a concise challenge prompt and solicit varied solutions. NYC has already used challenge procurements to find novel cooling solutions for public housing and to seek partners for a new tech campus. *Making better use of private-sector knowledge just makes sense* in a fast-moving tech landscape and challenge-based solicitations can unlock that knowledge. Miami-Dade's successful use of pilot-based open innovation challenges model shows it can cut procurement timelines by over 67% and increase access and innovation[87].

Problem-first procurements are especially suited for emerging tech like generative AI where agencies may not know the best solution up front. Instead of writing rigid requirements for a specific AI tool (which could be outdated by award time), DOTs can describe their transportation or operational challenge (e.g. optimizing traffic flows, automating data analysis, improving customer service via AI) and let vendors propose creative AI solutions. This approach reduces procurement cycle time (no lengthy spec development) and increases vendor diversity and innovation. It also allows agencies to evaluate real prototypes – ensuring the chosen AI solution actually works on their problem – rather than selecting based on paper proposals alone. Going forward, we can expect challenge-based RFPs (like California's RFI² model) to become a standard tool for AI procurements in government.

- **Modular and Phased Contracts**: For large-scale systems or long-term needs, DOTs are moving away from big bang contracts toward modular procurement and phased implementations. In a modular contracting approach, a major project is broken into smaller, bite-sized modules that can be bid out separately and delivered incrementally[88]. For example, instead of one huge contract to build an intelligent transportation system end-to-end, a DOT could issue separate modules (data ingestion, analytics engine, user interface, etc.) potentially to different specialized vendors. This has several benefits: it introduces flexibility (if one module fails or a technology changes, you can course-correct without derailing the entire project), lowers risk and cost of failure, and allows smaller innovative firms to compete on pieces of the project rather than requiring one giant prime contractor.

  Phased delivery is also related – the contract is structured in stages with go/no-go decision points. A DOT might contract for Phase 1 to develop a prototype or pilot, then only proceed to Phase 2 (full deployment) if Phase 1 meets defined success criteria. This approach, often combined with challenge-based procurements, ensures that unproven AI solutions must *earn* the larger rollout by first demonstrating value. It also builds agility into contracts: agencies can incorporate new requirements in later phases as technology evolves. These Modern RFP processes also encourage innovative solutions to identify the most qualified bidders[89]. For example, using a prototype challenge, where vendors build and demonstrate a working model to prove their capabilities. The

[85] https://media.governmentnavigator.com/media/bid/1704502042_1_4_2024-0000029442.pdf
[86] https://www.route-fifty.com/finance/2024/05/challenge-based-procurement-just-makes-sense/396942/
[87] https://medium.com/miami-dade-innovation-authority/its-time-to-reimagine-the-rfp-7b4964d866bc
[88] https://cdt.ca.gov/newsroom/2022/08/california-redefines-state-technology-procurement/
[89] https://cdt.ca.gov/newsroom/2022/08/california-redefines-state-technology-procurement/

down-selection minimizes the burden of evaluating dozens of full proposals and allows more interaction or negotiation with finalists. For AI projects, phased contracts mean you're not locked into a stale solution; you can pilot a generative AI tool in one department, then expand agency-wide only after confirming it's effective (or pivot if it's not).

- **Flexible Vendor Pools and Cooperative Purchasing:** Another strategy to accelerate procurement is to pre-position vendors and contracts so they're ready when needed. One method is establishing pre-qualified vendor pools through an open, ongoing process. California's Department of Technology, for example, created an Electronic Vendor Application of Qualifications (eVAQ) process – a continuous application where vendors submit their qualifications and can be added to a pre-qualified list at any time[90]. When a need arises (say a new AI tool or IT service), the agency can invite those pre-vetted vendors to bid. The eVAQ process streamlines the submission and assessment phase for repetitive administrative requirements.

  Many states also maintain standing contracts or master service agreements for IT services that agencies can leverage instead of starting from scratch[91] [92]. If a state has a master contract for AI cloud services or is part of a cooperative purchasing agreement, a DOT could issue a task order or purchase under that umbrella rather than a full RFP. For instance, the California Department of General Services, Procurement Division (DGS-PD), offers pre-negotiated contracts through Leveraged Procurement Agreements (LPAs), allowing public entities like state and local governments, schools, and universities to purchase goods and services without further competitive bidding under Public Contract Code Section 10298[93]. Using cooperative contracts (inter-state or inter-agency agreements) is explicitly recommended as a way to streamline procurement and avoid reinventing the wheel. It gives access to competitively-bid contracts that other governments have already put in place, thereby cutting months off the schedule for new procurements.

  In summary, having flexible procurement vehicles – whether state-run vendor pools or cooperative contracts – allows government buyers to move quickly when a new AI solution is needed, because the legal contracting framework is already in place. This is particularly useful for lightweight tools or departmental pilots: for example, a DOT traffic operations division could rapidly procure a subscription to an AI-driven traffic modeling software if the vendor is pre-approved on a state IT contract.

- **Embracing Pilot Projects (and Avoiding Pilot Purgatory):** Experimentation is vital with emerging tech, so agencies are encouraging small pilot projects to test AI capabilities. The challenge has been how to integrate pilots into procurement without getting stuck. Often, governments run an isolated pilot with a startup or new tool, but when it succeeds, procurement rules force them to do a whole new RFP to scale it – losing momentum and frustrating both agency and vendor. New approaches aim to eliminate this *pilot purgatory*. The idea is to design the procurement process to include the pilot and the path to scale in one package. As noted above, California's RFI² and similar challenge-based processes do exactly this: they fund a pilot with one or multiple vendors under the initial solicitation, then allow the winning pilot to convert into the full implementation contract if successful[94]. DOTs can similarly set up pilots with clear success metrics and an *option* to extend or expand the contract if those metrics are met. This could be a contract with phases (pilot phase and expansion phase as noted). The bottom line is to institutionalize fast transitions

[90] https://cdt.ca.gov/project-delivery/statewide-technology-procurement/
[91] https://www.dgs.ca.gov/PD/Resources/SCM/TOC/16/16-10
[92] https://www.dgs.ca.gov/PD/Resources/Page-Content/Procurement-Division-Resources-List-Folder/IT-Consulting-MSA-5167010
[93] https://www.dgs.ca.gov/PD/About/Page-Content/PD-Branch-Intro-Accordion-List/Acquisitions/Leveraged-Procurement-Agreements
[94] https://www.route-fifty.com/finance/2024/05/challenge-based-procurement-just-makes-sense/396942/

from pilot to production – otherwise agencies remain stuck perpetually testing technology without fully adopting it.

Additionally, some states are creating innovation sandbox programs for AI. California's recent executive order on GenAI directs the state to create infrastructure to allow state agencies to carry out pilot projects and sandbox testing of new GenAI applications[95]. While details are forthcoming, this suggests a framework where agencies (like a DOT) could experiment with AI tools in a controlled environment (with appropriate oversight on data, ethics, etc.).

Regulatory sandbox models let innovators test under relaxed rules for a period of time. Applied to procurement, a sandbox might, for example, permit a DOT to do a time-bound trial of an AI system with simplified procurement procedures or exemptions, so long as it stays below a certain scale and duration. Texas' new AI governance law[96] contemplates allowing businesses to test AI systems for up to 36 months in a controlled setting without full regulatory compliance. For government agencies, a similar pilot/sandbox authority could be invaluable – enabling quick, low-risk trials of AI solutions. The recommendation for DOTs is to advocate for or utilize pilot-friendly authorities: whether through explicit state programs, special procurement designations (e.g. treating something as an R&D experiment), or simply using the smallest viable procurement method (like purchase card or micro-purchase[97]) to trial a new AI tool on a limited basis.

- **Technology & Training to Streamline the Process**: Part of reimagining procurement is modernizing the process itself with better tools and training – effectively *procurement transformation*. Many procurement offices are upgrading from traditional, paper-based or siloed systems to modern e-procurement and document automation platforms[98] [99]. government procurement can be a powerful tool for innovation and local economic growth if modernized. procurement offices must move beyond outdated, rigid systems by adopting agile, user-centered approaches that simplify bidding, increase competition, and improve access. Key strategies include streamlining RFPs, using cooperative purchasing, tracking outcomes, and integrating procurement into broader policy goals. When paired with investments in staff training and digital tools, procurement can become a strategic driver of efficiency, equity, and innovation in government operations[100].

  Leveraging AI tools to speed up procurement workflows is already yielding results in some governments. State procurement officials have been exploring generative AI to draft RFP documents, summarize vendor responses, and analyze contracts much faster than before. Early findings across states show AI-driven assistance can shrink portions of the process from months to weeks, weeks to days, or days to hours. In one case, Idaho used an AI tool in its IT procurement reviews and was able to compress a 6-12 month vendor evaluation process into a much shorter timeframe by automating security assessments[101]. Likewise, cities like Boston are developing AI chatbots trained on procurement regulations to help staff quickly navigate rules and get purchases approved faster[102].

  Beyond tools, building a procurement innovation culture is key. California's Department of Technology, for instance, set up a Technology Procurement Community of Practice to share best

[95] https://www.gov.ca.gov/wp-content/uploads/2023/09/AI-EO-No.12-_-GGN-Signed.pdf
[96] https://www.nelsonmullins.com/insights/alerts/privacy_and_data_security_alert/all/texas-legislature-passes-house-bill-149-to-regulate-ai-use
[97] https://factoringclub.com/government-contract-factoring/government-micro-purchase/
[98] https://authorium.com/exploring-the-latest-trends-in-government-procurement/
[99] https://insider.govtech.com/california/sponsored/the-two-procurement-concepts-that-will-enable-california-to-make-the-most-of-taxpayer-dollars
[100] https://manhattan.institute/article/the-power-of-the-purse-how-to-make-government-procurement-a-tool-for-innovation-and-growth-in-local-economies
[101] https://statescoop.com/generative-ai-government-procurement-risk-2024/
[102] https://bloombergcities.jhu.edu/news/strategies-spreading-ai-throughout-local-government

practices and train procurement professionals in agile methods[103]. The aim is to create a network of experts who understand modern IT procurement, since procurements are complex and require substantive expertise in the specific domain (like AI). Similarly, the DHS Procurement Innovation Lab not only runs pilot projects but also conducts bootcamps to train thousands of procurement officials in new techniques[104]. DOTs should ensure their procurement staff (and project managers) are educated in these emerging methods – whether it's how to run a challenge competition, how to evaluate an AI solution's risk, or how to negotiate flexible contract terms. This human capital element will help institutionalize faster and smarter procurement long-term.

Recommendations for Future DOT AI Procurements

Considering the above review, here are actionable recommendations for state DOTs and other agencies to modernize procurement in the era of AI. These apply to both enterprise-level projects and smaller departmental pilots:

- **Adopt a Problem-First Approach**: When seeking AI solutions, start with a clear *problem statement* or outcome goal in your solicitation. Encourage vendors to propose innovative AI approaches rather than prescribing every detail. This challenge-based procurement model (e.g., California's RFI²) attracts a wider range of ideas and can significantly shorten RFP documents while yielding higher-quality solutions.
- **Build Phases into the Procurement**: Structure solicitations to include an initial pilot or prototyping phase with a smaller award, followed by an expansion phase for the successful solution. Make it explicit that if the pilot meets defined metrics, the vendor can swiftly proceed to implementation without a full rebid. This not only motivates vendors to perform well, but also avoids losing 6-12 months in pilot purgatory. Ensure the legal team crafts the contract or solicitation to allow this (e.g. use contract options or multi-stage awards as permitted by the state's laws).
- **Leverage Flexible Procurement Vehicles**: Don't always start from zero. Use existing cooperative contracts, state master agreements, or pre-qualified vendor pools to source AI tools and services. For instance, if another state or a federal GSA schedule has an AI software contract, explore if your DOT can piggyback on it to get a solution in weeks. Internally, consider establishing a continuous qualification program for emerging tech vendors (similar to eVAQ) so that you have a go-to roster of vetted AI solution providers ready for quick bids.
- **Shorten Procurement Cycle Time** by modularizing big projects and using agile RFP processes. For large, mission-critical systems (e.g. an intelligent infrastructure management platform), break the procurement into modules or phases. This allows technology updates or vendor swaps module- by-module if needed, preventing lock-in to outdated tech. Also, use *Modern RFP* practices: shorter initial proposals, a down-select to a few finalists, then deeper engagement with those finalists including possible solution demos or interviews. This way, evaluation is faster and more interactive, and you select the most qualified, innovative partner instead of the best proposal writer.
- **Institute Agile Contracts**: Wherever possible, craft contracts that are flexible to change. This could mean including technology refresh clauses (requiring the vendor to periodically update the AI model to state-of-art versions), allowing adjustments via change orders or addenda as new AI capabilities emerge, or using performance-based requirements (focus on accuracy, speed, outcomes) rather than fixed technical specs. An agile contract combined with agile project management means the DOT can continuously incorporate improvements during the contract term, crucial for AI where

[103] https://cdt.ca.gov/newsroom/2022/08/california-redefines-state-technology-procurement/

[104] https://www.deloitte.com/content/dam/assets-zone1/tw/tc/docs/about/2025/government-trends-2024-en.pdf

algorithms rapidly evolve.

- **Empower Departments to Pilot**: Enable your program-level or district-level teams to experiment with AI tools on a small scale without excessive bureaucracy. Set guidelines for safe harbor pilot projects – e.g. under a certain dollar threshold or using non-production data – that can be executed with streamlined approval. Many states allow small purchases under X dollars with simplified procedures; use that to let a traffic ops team, for example, quickly try an AI-based routing tool. Additionally, push for state-level initiatives (or executive orders) that support GenAI sandbox pilots. If your state offers an innovation grant or has R&D authority, tap those channels to run proofs- of-concept. The goal is to learn by doing on a small scale, then leverage those lessons for a formal scaled procurement.
- **Improve Internal Processes with AI**: As you buy AI, also *use* AI to buy better. Encourage procurement offices to adopt tools like AI-driven document analysis, chatbots, or workflow automation to accelerate their work. This can cut down the administrative lead time significantly – for example, using generative AI to draft an RFP or contract can shave weeks off drafting and ensure compliance by referencing past solicitations. Training an AI assistant on your DOT's procurement rules could help project managers quickly navigate compliance. By automating repetitive tasks in solicitation development, bid evaluations, and vendor communications, procurement professionals can focus on high-value judgments and market research. Generative AI can potentially reduce parts of the process from months to days, but it requires proper risk management and validation. DOTs can pilot these internal AI tools within their procurement divisions (with necessary oversight) to gain speed while maintaining fairness and transparency.
- **Invest in Procurement Workforce and Policy Changes**: None of these innovations happen without updating policies and mindsets. DOT leadership can work with state CPOs and legislators to modernize procurement regulations, allowing for things like challenge-based RFPs or negotiations in the RFP process if those aren't currently permitted (e.g., an executive order to authorize RFI² competitions) . Push for the needed statutory flexibility to procure rapidly evolving technology in non-traditional ways. At the same time, invest in training procurement staff on new practices (whether through formal certification in agile procurement or informal communities of practice). Celebrate and replicate successes (for instance, if a pilot succeeded via an innovative contract, use it as a case study). A cultural shift is needed so that procurement is seen not as a hurdle but as a strategic partner in innovation. When everyone – from engineers to contracting officers – understands agile, outcome-focused procurement, the agency can move with much greater speed and confidence.

Summary of Findings

GenAI has emerged as a transformative technology with significant potential for improving government operations, including traffic management, customer service, and data analysis. However, it presents considerable risks such as algorithmic bias, security vulnerabilities, and accuracy challenges. Federal and state-level procurement practices for GenAI now integrate traditional IT procurement guidelines with additional AI-specific standards, emphasizing risk management, data privacy, transparency, and ongoing oversight.

Procurement policy frameworks, such as those from the OMB and GSA, stress outcomes-based procurement, the inclusion of privacy officials, cross-functional governance teams, and early piloting in controlled environments. State initiatives, notably California and Washington, have developed detailed procurement guidelines aligned with NIST's AI Risk Management Framework, prioritizing fairness, equity, and public-interest alignment.

Vendor selection now requires evaluations focused on fairness, algorithmic transparency, security compliance, and performance metrics. Contracts are increasingly structured to avoid vendor lock-in and protect government data and intellectual property. Additionally, data privacy standards mandate strict restrictions on using government data for AI training, and cybersecurity compliance aligns with existing frameworks such as FedRAMP.

Implementation mechanisms include forming interdisciplinary teams, launching pilot programs, using structured procurement methods, offering continuous training, enforcing security and compliance standards, and setting up ongoing monitoring and evaluation processes. For Caltrans, the next steps involve developing a GenAI procurement framework, engaging key internal stakeholders, identifying initial pilot use cases, setting up governance structures, building capacity through training, and establishing performance metrics aligned with operational and ethical goals.

state DOTs must reimagine their procurement playbook to keep pace with AI advancements. This means embracing agility at every stage: from identifying needs as problems to solve (not fixed solutions to buy), to running phased, challenge-based competitions that produce working AI prototypes in months, and to leveraging modular contracts and vendor pools for fast scalability. It also means creating pathways for both large-scale systems and lightweight departmental tools – the former through flexible, modular multi-year procurements, the latter through quick pilots and sandbox experiments that can graduate to broader deployments. These practices, already being pioneered in pockets of government, are likely to become the new standard. By adopting the approaches above, DOTs and other agencies can shorten procurement cycles dramatically without sacrificing fairness or rigor. In fact, they may improve outcomes by inviting more innovation and competition into the process. The payoff will be the ability to harness generative AI and other emerging tech when it's needed, not years after the fact – ultimately delivering better transportation solutions to the public in real time.