PHYSICAL SECURITY | CYBERSECURITY | OPERATIONS

AI IN ELECTIONS

Procurement Essentials

SAFEGUARDING ELECTIONS THROUGH RESPONSIBLE ACQUISITION.

ELECTION SECURITY *Exchange*

The integration of Artificial Intelligence (AI) into elections represents a fundamental transformation in how organizations operate and make decisions. Unlike traditional information technology, where functions yield predictable outputs, AI systems are based on probability and adaptive.

This shift necessitates reimagining the procurement lifecycle because AI poses unique risks that traditional procurement may overlook.

As a follow-up to our series introduction, AI in Elections: Essentials for Risk Management, this guide offers practical advice for the responsible adoption and procurement of AI-enabled tools, especially for election systems that affect individual rights and safety and require the highest level of protection for managed data, such as voter records and infrastructure details.

In this guide, you will learn:

- **Pre-Procurement Planning:** How to establish your use case and the necessary foundation before engaging with vendors.
- **Drafting the Solicitation:** Understand how to integrate standard IT best practices with AI-specific guardrails.
- **Vetting:** Discover why human verification is a non-negotiable requirement and questions to ask of vendors.
- **Post-Award Monitoring:** Learn why AI oversight does not stop at contracting signing, and how to build in ongoing review.

Phase 1: Pre-Procurement Planning

Successful acquisition begins long before a written solicitation is issued. Before you even start drafting, you need to determine whether AI is actually necessary and safe for your specific use case.

1. Define Your Use Case

Before evaluating any AI solution, just like any other technology solution, you need to articulate the problem you're attempting to solve clearly. Do not acquire or use AI simply because it's the latest trend. What specific administrative pain point are you trying to solve? The best way to eliminate potential risk vectors from AI is not to introduce them if there is no legitimate business need to do so.

- Clarity check: When writing your requirements, be specific without prescribing the solution.
  - **Unclear:** "We need a tool to improve communication with voters."
  - **Better:** "We need a system to reduce average response time to inquiries from 24 to 12 hours while maintaining 95% accuracy."

Bottom line: Never procure AI to look for a problem. Define your use case first, then verify if AI is an effective option.

2. Assemble the Dream Team

Given the complexity of AI, which involves legal, security, and public perception considerations, vetting potential solutions on your own is difficult.

To address this, form a team with the necessary expertise and perspectives to guide the strategy and requirements *before* drafting the solicitation. The composition of this team should be determined well in advance of beginning any formal procurement process.

- For example, **the election official** may define the problem, the chief information officer/ chief information security officer may vet the system architecture, **legal counsel** may review liability, and **an external partner, such as a university or cybersecurity expert**, may review the vendor's technical claims.

3. Audit your Policies

In elections, audits are ubiquitous. Before seeking new AI tools, organizations must look inward. Standard IT policies often focus on cyber hygiene and infrastructure security, but don't fully address the specific complexities of AI models. Therefore, this is an ideal time to establish or update your AI use policies and data governance structure.

Implementing this framework is essential, as it will properly guide your solicitation process and define necessary vendor requirements. Microsoft has assembled an AI Use Policy Template for election offices as a practical starting point, covering key components, including definitions, roles, and responsibilities. For a more comprehensive governance framework, including a structured review process, the GovAI Coalition AI Governance Handbook is designed for public agencies and can be adapted for any sized election office.

- For example, when it comes to data governance, you'll want to consider how the vendor can or can not use your data, including where the data should be stored. CISA, the National Security Agency, the Federal Bureau of Investigation, and international partners released a joint Cybersecurity Information Sheet on AI Data Security that provides more in-depth guidance on securing training data and operating AI-enabled systems.

4. Plan for Humans

Particularly in election administration, AI should never be the final decision-maker. The Brennan Center for Justice notes that human involvement and review can help mitigate concerning behavior and bias. Similar to bipartisan review teams that review ballots, jurisdictions should implement rigorous requirements at the outset to ensure that humans are involved at every step of the process so that systems remain accurate and effective without causing harm.

5. Conduct Market Research

Election officials' procurement strategy for AI should also consider an analysis of the market and available vendors. Unlike traditional IT procurement, which often focuses on financial stability and past performance, assessing and rating potential AI vendors must encompass a broader range of criteria. These include model provenance, ethical data sourcing, and the data supply chain. If this information is not readily available, it should be explicitly requested in the solicitation documents. Market research can also surface whether variations of AI tools exist for government use. Typically, these versions offer greater data sovereignty and security, improved compliance with security benchmarks, such as FedRAMP certification, and options to segregate or limit servers from the public internet where AI tools are enabled.

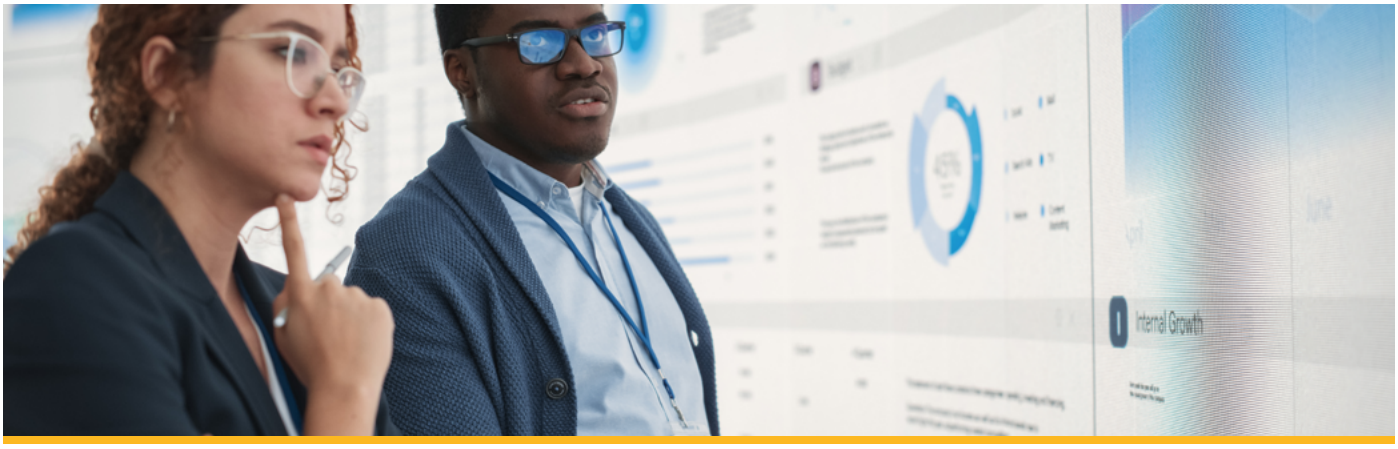

Phase 2: Drafting the Solicitation

1. Know the Landscape

AI-enabled systems differ fundamentally from traditional software. To evaluate them effectively, you must explore them through two lenses: standard IT best practices and AI-specific guardrails. As noted by the National Association of State Procurement Officials (NASPO), cybersecurity should be integrated throughout the process and be considered from the beginning.

- **Standard IT Procurement (Baseline):**

  Examples of standard procurement practices, particularly in the public sector, that you should already apply to any technology solution include:

  - **Alignment:** Ensure procurement procedures and standards are aligned with your organization's policies. The State of North Carolina has established in-depth reporting requirements for assessing state-hosted and non-state-hosted solutions.
  - **Access Control:** Ensure the system supports role-based access and multifactor authentication (MFA).
  - **Uptime & Availability:** Traditional software focuses on reliability. Typically, service level agreements (SLAs) help enforce or guarantee a specific level of uptime during key business hours or election periods.
  - **Data Encryption:** Mandate encryption for data in transit and at rest, using approved standards such as AES-256.
  - **Validate Cyber Hygiene:** Understand the vendor's cyber risk management practices, including how they identify and manage vulnerabilities and how they routinely apply patches.

  For more comprehensive best practices and model procurement language, consider the Center for Internet Security's (CIS) Guide to Election Technology Procurements.

- **AI-Specific Procurement:**

  AI introduces more dynamic risk that may fall outside the scope of traditional procurement.

  - **What's Under the Hood?** CAI (https://www.cai.io/) outlines the importance of transparency in reasoning in their white paper on responsible AI guidelines. Understanding decision-making is crucial for strengthening trust, validating output, and enhancing fairness and accuracy. Without it, the system can be a "black box," preventing its actions from being explained.

    RECOMMENDATION: *Require the vendor to document the model's decision-making process and lifecycle in detail.*

  - **Data Training:** With standard software, you own the data. With AI, it can be different, where your data is used to train the base system.

RECOMMENDATION: *Inquire about data sources used to train the model(s), and what mechanisms are in place to protect your organization's data.*

- **Thoughts can "Drift":** With AI systems, responses can shift over time based on new information, queries, and responses generated during use, and on user feedback on the quality of generated content. Be mindful of the impact of these types of changes and how they may affect the usability and reliability of responses.

  RECOMMENDATION: *Require robust monitoring and periodic evaluation to verify the tool's accuracy and outputs, and always require human verification before producing a final result.*

- **Privacy Concerns:** Often, according to the National Institute of Standards and Technology (NIST), privacy clauses fail to account for AI's ability to infer sensitive information from "non-sensitive" data. AI systems can present this risk, even if the data was anonymized.

  RECOMMENDATION: *Employ the principle of data minimization, only providing the data that is absolutely needed to function. Also, require vendors to use Privacy-Enhancing Technologies (PETs).*

- **Factoring in Fairness:** AI models can amplify existing biases. NIST notes that systems might be "accurate," but they disproportionately fail for particular demographics.

  RECOMMENDATION: *Do not accept accuracy metrics at face value. Require "disaggregated evaluation" that breaks out performance across demographics. Human oversight can also help mitigate bias before it impacts a voter.*

2. Embrace the "Clause"

When drafting a solicitation of this type, it is important to clearly define the guardrails for AI and inquire about infrastructure and data model specifications. This is where the multidisciplinary team "Dream Team" you've formed can shine and add significant value during the drafting process. A few key areas to consider, with example language to help you get started if you need it, are listed below:

- **Explainability:** An example may look like this: "The vendor must provide a plain-language explanation of how the system reaches its conclusions. Algorithms that can not be audited for reasoning are not acceptable."
- **Training Data Transparency:** A possible example of a requirement is: "The vendor is required to disclose the data sources used to train the model, specifying whether they include public voter data, proprietary data, or data from the open web. Furthermore, the vendor must warrant that they possess the legal rights to all training data used."
- **Bias Testing:** For example, "The vendor must provide documentation of pre-release bias testing, specifically regarding demographic disparities (e.g., race, age, etc) in performance accuracy."

- **Data Ownership:** For example, “The jurisdiction retains exclusive ownership of all input data. The vendor is strictly prohibited from using the jurisdiction’s data to train their base models for other clients.”
- **Hallucination Risk:** For generative AI tools, you might consider something like: “The vendor must detail the technical guardrails in place to prevent the generation of false or misleading information (hallucinations).”
- **Incident Notification:** Timely vendor disclosure of security or performance issues is paramount. For example, “The vendor shall notify the jurisdiction within 24 hours of discovering any security breach, unauthorized access, or material change in model accuracy.” Time-based requirements are a significant factor in product costs, regardless of whether the product or service is AI-enabled. Therefore, each jurisdiction must assess the product's or service's purpose and align it with its organization's acceptable risk level.
- **Transition:** Another item for consideration is protecting the jurisdiction’s ability to move away from a procured system or set of tools. For example, “Upon contract termination, the vendor shall return all jurisdiction data in a standard, machine-readable format within 30 days and provide written certification that all copies of jurisdiction data have been permanently deleted from vendor systems, including data contained in backups. The vendor shall provide reasonable transition assistance, including documentation and data for a period not to exceed 90 days post-termination.”

The type of desired tool will drive the specific requirements and questions to ask potential vendors. The European Commission’s Public Buyers Community Platform offers model contract clauses for AI procurement based on the European Union’s Artificial Intelligence Act (EU AI Act). Washington State’s Enterprise Services also provides guidance on AI contract clauses.

3. Define Performance and Accuracy

When translating your needs into requirements, officials must distinguish between traditional IT metrics and AI-specific outcomes. This often requires ongoing monitoring and accountability mechanisms, such as audit rights, as noted in the State of Georgia’s enterprise policy on the Procurement of AI Tools Guidelines for Responsible Use. For example, it’s important to incorporate outcome-based requirements when working with adaptive systems.

- **Functional vs Outcome Requirements:** A more traditional requirement might be purely functional: “The system must process 500 forms per hour.” For AI-enabled systems, this is not sufficient. Organizations will need to define the outcome and tolerance for error, such as: “The system must extract data with 95% accuracy and maintain a false positive rate of less than 1% for fraud detection.”
- **Consider Fairness:** Depending on the use case, it may be necessary to loop in fairness criteria. For example, organizations may want to state explicit metrics the vendor must meet and require performance reports disaggregated by specific categories (e.g., race, gender, and age) to reveal hidden bias.

Phase 3: Vetting

You can't just rely on a slide deck. You'll need to verify and validate all claims and responses. It's also a best practice to invite respective bidders to provide a demo, which is an excellent opportunity to ask more specific questions.

1. *Q&A with Vendors*

Use the demonstration to verify the vendor's claims by asking detailed questions about their processes and expertise. If a vendor is unable to respond with clarity or precision, this suggests an area that warrants further scrutiny.

- "When a model makes a mistake, what is the remediation?"
- "Can you show us how you're aligned with NIST's AI Risk Management Framework (AI RMF)?
- "What data is used to train the model, and who labels it? Is it done in-house or outsourced?"
- "What is your model update and versioning policy?"
- "How do you notify customers of any changes to the model that could affect accuracy, fairness, or behavior?"
- "Do you have a published AI incident response plan?"

2. *Security Qualifications*

At the end of the day, AI tools are software that rely on hardware, just like voter registration databases or other election systems, though AI tools are much more complex. The vendor should be able to demonstrate adherence to rigorous standards critical to election infrastructure.

Don't be afraid to ask the vendor to provide proof of industry certifications (e.g., FedRAMP or StateRAMP authorization) or a recent penetration test report that targets AI attacks, such as tricking the chatbot into giving incorrect information or extracting sensitive information about the model. For organizations looking to leverage generative AI, the Open Worldwide Application Security Project (OWASP) lays out vulnerabilities and mitigations in the OWASP Top 10 for LLM Applications 2025 to strengthen security and engagement with vendors.

Additionally, ask vendors to provide a Software Bill of Materials (SBOM) or an AI Bill of Materials (AIBOM) that identifies all third-party components, pre-trained models, and external data sources incorporated into the system. The DHS Roles and Responsibilities Framework for AI in Critical Infrastructure identifies supply chain documentation as a key obligation for critical infrastructure.

3. *Data Provenance and Governance*

Since election infrastructure is highly sensitive, data management is a must. Ensure that the data provided to the AI vendor is secure and that its origin is tracked. The vendor must be able to prove that it cannot expose or misuse sensitive information, especially if the use case requires sensitive voter information. Mechanisms such as cryptographic signatures, trusted infrastructure leveraging Zero Trust principles, and decentralized learning techniques can be helpful, as noted in the Cybersecurity Information Sheet referenced earlier.

Phase 4: Post-Award Monitoring

The contract signing is merely the beginning of the responsible deployment of AI technology, not the end. Unlike static, traditional software, AI systems are dynamic; their behavior can shift as they process new data, models are updated, or deployment environments change.

The thorough work completed in Phases 1 through 3 creates a necessary foundation. However, it is the accountability embedded in the ongoing vendor relationship that will determine the long-term stability of this foundation. The NIST AI Risk Management Framework clearly establishes post-deployment monitoring as a core requirement rather than an optional step. For election officials, who bear the ultimate responsibility to their constituents, this phase is when their commitment to stewardship is most evident.

1. Establish Ongoing Performance Monitoring

Continuous, structured oversight, beyond a single acceptance test, is necessary for AI systems. Ongoing monitoring and audit rights are vital for accountable AI adoption, as highlighted by the State of Georgia's enterprise policy on Procurement of AI Tools Guidelines for Responsible Use.

Performance must be routinely re-evaluated throughout the contract lifecycle, not solely at renewal. It should be judged against the original outcome-based requirements (see Phase 2). This is crucial because accuracy can degrade when real-world data deviates from training data, a phenomenon known as "drift". Similarly, fairness outcomes may shift with the introduction of new demographic groups or edge cases into the workflow.

- Establish a performance review schedule for systems: at least quarterly for high-impact systems and mandatory reviews both immediately preceding and following every election cycle.
- Require the vendor to furnish regular performance reports. These reports must be disaggregated by relevant demographics and use cases, aligning with the fairness metrics defined when the contract was awarded.
- Ensure the contract includes the right to conduct or commission independent audits of the system's performance, data handling, and security controls throughout the entire contract term.

2. Prepare for AI Incidents

AI incidents encompass more than just cyberattacks; they include accuracy failures, bias events, hallucinations affecting voters, and the inadvertent use of jurisdictional data to train a vendor's base model.

Every jurisdiction using AI must have a proactive plan to detect, respond to, and recover from these events. A practical starting point for election offices of any size is the GovAI Coalition AI Incident Response Plan. This ready-to-use template is built upon the NIST AI Risk Management Framework and the NIST SP 800-61 Incident Response Recommendations and Considerations for Cybersecurity Risk Management.

- Establish an AI Incident Response Team before deploying any AI system. Define roles, assign ownership, and conduct tabletop exercises at least annually to test your response procedures.
- Know your activation criteria. Triggers should include declining accuracy rates, suspected bias events, failed human oversight controls, or any discovery of unauthorized use of jurisdiction data.
- Ensure critical operations have a documented backup plan that does not rely on the AI system. If the tool goes offline or is suspended mid-cycle, your office must continue functioning.

3. Manage Updates and Changes

Model updates from vendors are a standard, often positive, occurrence. However, for election officials, an unannounced update that modifies how an AI system processes data or generates recommendations constitutes a material change. Election officials should manage model updates with the same rigor applied to software changes in any other critical system. This means requiring advance notification from the vendor, establishing a formal approval or opt-out process, and revalidating the system's performance after any significant update before deploying it back into production.

- Require contractual advance notice of any material model update (a minimum of 30 days is a reasonable baseline), along with a summary of what changed and what testing the vendor conducted.
- Maintain a change log. Document each model version your jurisdiction uses, when it was deployed, and the results of your post-update performance validation. This creates an auditable record and supports accountability if a problem is later traced to a specific update.

4. Know When to Decommission

Not every AI system should remain in service indefinitely. The NIST AI RMF is direct on this point: Organizations should decommission systems that exceed established risk tolerances. Election officials should proactively define the exact performance thresholds or conditions, ideally as part of the contract, that would necessitate suspending or completely removing an AI system.

Clearly documenting these criteria eliminates uncertainty, allowing officials to take quick, decisive action if the system fails to meet expectations or presents a risk to the electorate.

- Define decommission triggers in your contract. Examples include: accuracy falling below a defined threshold for two consecutive reporting periods, a confirmed bias event affecting a protected class, a material security incident, or a vendor failure to remediate a known defect within the agreed timeline.
- Invoke the transition provisions established in Phase 2. A well-drafted exit clause means that decommissioning a problematic tool does not leave your jurisdiction without data access or operational continuity.

Conclusion

Procuring AI systems is not just about buying software; it's an exercise in risk management and in stewarding public trust. In an interconnected world, security is not optional. Election officials who ground their procurement decisions in sound principles are better positioned to address the dynamic risks that AI systems introduce. The choices made throughout this process, whether contract clauses, transparency requirements, or human oversight mechanisms, will mold the integrity of the solution and its role in election operations.

By anchoring procurement in these principles and strategies, election officials can strengthen the responsible use of AI. The goal is not simply to acquire efficiency or capacity; it is to build and sustain public confidence in election operations. Proceeding with care and with AI-specific security considerations as a guide helps ensure that operations remain secure and trusted.

Procuring AI systems is not just about buying software; it's an exercise in risk management and in stewarding public trust.

Further Resources and Templates

Additional resources and templates are provided below to supplement the information in this document. This is not an exhaustive list, but it provides a helpful starting point. Election officials should review these resources and share them with key partners to prepare for AI adoption and mitigate associated risks.

Simple and Quick

- AI in Elections: Essentials for Risk Management
- AI Use Policy Template for Election Officials
- Government-Specific AI Tools: What Election Offices Should Know
- Safeguards for Using Artificial Intelligence in Election Administration

Technical and Detailed

- AI Risk Management Framework (AI RMF 1.0)
- NIST AI RMF Playbook
- Roles and Responsibilities Framework for AI in Critical Infrastructure
- Cybersecurity Information Sheet: AI Data Security
- OWASP Top 10 for LLM Applications 2025
- Buyer Be Aware: Integrating Cybersecurity into the Acquisition Process
- Guide to Election Technology Procurements
- Responsible AI Guidelines: A Blueprint for Ethical Integration
- FedRAMP Marketplace | General Services Administration
- Vendor Readiness Assessment Report
- Incident Response Recommendations and Considerations for Cybersecurity Risk Management (SP 800-61r3)

Templates

- AI Governance Handbook
- AI Incident Response Plan Template
- Procurement of AI Tools: Guidelines for Responsible Use
- Model Contract Clauses for AI Procurement
- AI Contract Clauses Guidance