



The Public Sector Guide to Generative AI

Introduction 3

Why Generative AI?..... 3

How to Get Started with Generative AI..... 4

Step 1: Establish Readiness..... 4

Step 2: Evaluate Generative AI Tools 6

Step 3: Get Stakeholder Buy-In7

Step 4: Define Success Upfront 8

Next Steps for Public Sector Leaders 8

Introduction

For public sector agencies, mission delivery depends on the ability to process vast amounts of information quickly, accurately, and transparently. Whether managing complex litigation, responding to Freedom of Information Act (FOIA) requests, conducting investigations, or crafting policy and rulemaking, government legal, records, and compliance teams are tasked with performing data-intensive processes under intense scrutiny. Generative AI offers new ways to support their mission, not by replacing staff expertise, but by extending capacity in review tasks where oversight and defensibility are essential.

This white paper explores the growing role of generative AI in the public sector and provides a practical framework for understanding, evaluating, and adopting generative AI solutions in a secure and defensible way. From building organizational support to defining what success looks like, this is a resource for agencies seeking to responsibly apply technology in a time of rapid change.

Why Generative AI?

For more than a decade, public sector teams have relied on traditional machine learning tools like technology-assisted review (TAR) to accelerate document review. But while TAR reduces the burden, it has limited flexibility and requires significant upfront training. Purpose-built generative AI tools offer a more agile and user-friendly solution, with large language models analyzing documents one at a time and explaining their results in plain language. They can back up each output with clear citations and adjust dynamically to the reviewer's instructions.

Relativity aiR for Review in RelativityOne Government demonstrates this advantage in practice. Built specifically for document review use cases, aiR for Review leverages the power of generative AI through Azure OpenAI Service in the FedRAMP-authorized Azure Government Cloud to deliver more efficient, consistent, and higher-quality results. With integrated workflows and natural language, it identifies impactful documents and provides clear rationales supported by citations, enabling teams to quickly and deeply understand their data and to uncover insights that might otherwise be missed. This combination of speed, accuracy, and transparency helps government practitioners accelerate review with confidence.

aiR for Review is also designed to be ready to use immediately, with no model training or custom coding required from alternative AI tools. Teams can provide the context of their review, whether a FOIA request, litigation matter, enforcement action, or investigation and easily leverage the power of large language models to accelerate their work while concentrating on the substance of their review. By combining specialization, defensibility, and deep-rooted security, aiR for Review lowers the barrier to entry for government use of generative AI.

How to Get Started with Generative AI

Step 1: Establish Readiness

Identify the Right Use Case

Successful adoption of generative AI in government begins with high-volume, high-burden tasks and repeatable processes. These are the areas where technology can demonstrate measurable value such as backlog reduction, improved consistency, and faster more consistent review, all while allowing government staff to stay in control. Starting in these areas helps agencies build confidence in the technology and lay the groundwork for broader adoption.

✓ **FOIA or Similar Records Requests:**

Handling FOIA requests requires labor-intensive document review, redaction, and response. Generative AI can help staff:

- Automate responsiveness review and flag potentially exempt content
- Accelerate turnaround time and reduce backlog
- Streamline initial triage of record sets
- Prioritize responsive records and focus on complex redaction decisions

✓ **Internal Investigations and Inspector General Reviews:**

Investigations often involve reviewing large volumes of emails, reports, and transcripts on highly specialized topics with sector-specific language or terms. With generative AI, investigators have the ability to:

- Surface relevant documents quickly based on context
- Summarize large sets of investigatory notes or interview transcripts
- Detect anomalies or patterns requiring additional scrutiny
- Prioritize high-risk documents early to focus human review where it matters most

✓ **Litigation and Regulatory Response:**

Legal and oversight responsibilities require agencies to manage significant volumes of evidence, subpoenaed materials, and regulatory filings. Generative AI can help agencies:

- Conduct first-pass review of discovery documents
- Highlight red flags or patterns of noncompliance for further review
- Surface arguments or evidence gaps across regulatory comments or submissions to ensure full transparency in reporting requirements
- Identify high-value evidence early, strengthening their case strategy while reducing discovery costs

✓ **Rulemaking and Public Comment Analysis:**

Generative AI can help agencies manage the flood of citizen feedback during rulemaking by organizing comments and ensuring they are reviewed fairly and transparently. It also helps policymakers:

- Identify comments responsive to specific provisions under review
- Group similar submissions for efficient handling
- Surface recurring themes or concerns across submissions to inform agency responses
- Accelerate the rulemaking and comment review processes while maintaining fairness and transparency

✓ Programmatic Mission Support:

Generative AI can help agencies reduce manual effort in routine program and administrative functions by classifying and organizing records for compliance and oversight. It enables agencies to:

- Identify records that may be subject to NARA retention schedules
- Surface documents requiring preservation for compliance, audits, or investigations
- Classify records by sensitivity or confidentiality level to guide handling
- Reduce manual review and burden of large-scale administrative record reviews

These use cases share common traits: they involve large volumes of unstructured information, require careful analysis, and depend on human oversight before final determinations are made. Together, these factors make them well-suited for applying generative AI to accelerate review while maintaining transparency and defensibility.

Build a Multidisciplinary Team

Generative AI adoption cannot be the sole responsibility of a single office. Success requires collaboration across the functions that safeguard compliance, manage information, secure technology, and deliver mission outcomes. Agencies that establish cross-functional task forces reduce blind spots and ensure adoption aligns with statutory requirements, oversight expectations, and mission priorities.

A well-positioned task force includes:

- IT/Security: Ensure secure and scalable integration
- Legal/Compliance: Safeguard statutory adherence and manage risk
- Program Leadership: Represent real-world workflow needs
- Management/Oversight: Provide accountability and alignment with mission priorities
- Budget/Finance: Assess resource requirements and long-term sustainability
- Procurement/Contracting: Guide acquisition strategy and ensure compliance with procurement rules

Choose a Pilot Project

Pilots give agencies the opportunity to demonstrate how generative AI can add value in practice. Rather than trying to apply generative AI across multiple workflows at once, focusing on single, high-impact use case – such as reducing FOIA backlog – can show measurable progress, build confidence with staff, and generate the results needed to inform future adoption and justify investment. A phased approach shows that generative AI can be deployed responsibly, with clear outcomes that strengthen trust in both the technology and the agency's oversight.

Step 2: Evaluate Generative AI Tools

Not all generative AI solutions are created equal. Introducing new technology in government requires confidence that the tools can be trusted to protect data, meet compliance obligations, and deliver defensible results.

Key evaluation criteria include:

Data Privacy: Generative AI models must ensure that agency data is not retained or used to retrain models outside agency control.

Fairness & Transparency: AI systems must provide explanations, rationale, and citations so reviewers can validate output and guard against hallucinations or errors.

Legal and Operational Specificity: General-purpose tools may be powerful, but without tuning for privilege, issue tagging, or responsiveness, outputs may be unreliable or indefensible.

These criteria informed our approach in building Relativity aiR for Review. We designed the solution to align with customer expectations for transparency, security, and defensibility. aiR for Review provides natural language explanations and citations in every output and also uses Azure OpenAI Service, which does not retain customer data after processing or use it to train Microsoft or OpenAI models.

Relativity's AI Principles

Generative AI offers government agencies new ways to reduce backlogs, accelerate review, and improve consistency across litigation, investigations, FOIA, and other record-intensive work. But introducing AI into public sector workflows also carries risks if it is not designed and deployed responsibly. For that reason, Relativity grounds its approach in six AI Principles that prioritize transparency, security, and accountability.

Relativity's AI Principles:

- 1 We build AI with purpose that delivers value for our customers.
- 2 We empower our customers with clarity and control.
- 3 We ensure fairness is front and center in our AI development.
- 4 We champion privacy throughout the AI product development lifecycle.
- 5 We place the security of our customers' data at the heart of everything we do.
- 6 We act with a high standard of accountability.

Step 3: Get Stakeholder Buy-In

Address Stakeholder Priorities

Generative AI adoption in government is not just a technical initiative – it’s an organizational one. Beyond pointing to efficiency gains, garnering support means addressing compliance concerns directly, showing how risks are managed, demonstrating accountability to oversight, and proving how generative AI strengthens the ability to meet your mission and better serve the public.

Functional Priorities for Generative AI

Stakeholder Group	Success Means ...	Focus Areas	What They Need to Hear/See
General Counsel	Reduced litigation burden; legally defensible outputs	Risk management, auditability, human oversight	AI outputs must withstand scrutiny in court, audits, or hearings; show audit logs and human review built in by default
CIO and IT Teams	Secure, sustainable integration of new technology	Integration, scalability, sustainability	Show how AI fits existing governance and cybersecurity frameworks without introducing new risks
Regulatory/Oversight	Timely rulemaking; transparent engagement	Fairness, transparency, independent validation	Demonstrate transparency features such as audit trails, clear documentation, validation opportunities
Program Offices	Capacity to focus staff on mission-critical work	Practical tools, training, clear workflows	Demonstrate how AI reduces manual workload; provide hands-on training and ensure clear feedback and escalation paths
Leadership	Mission delivery with accountability and transparency	Public trust, policy impact, responsible use	Position AI as delivering visible results while ensuring compliance, neutrality, and oversight alignment
Finance and Budget	Responsible stewardship of taxpayer dollars; efficient resource allocation	ROI, fiscal responsibility, sustainable spend	Show measurable savings, avoided contractor costs, and sustainable spend over the lifecycle

Step 4: Define Success Upfront

Success in generative AI is best measured across multiple dimensions – operational, quality, resource, and strategic metrics – to enable objective evaluation, justify investments, and guide future use. When defining success, consider:

Operational Metrics

- Time saved in document review processes
- Reduction in reviewer hours required
- Decrease in case or request backlog
- Volume of records reviewed within set timelines

Quality Metrics

- Precision and recall compared to human-only review
- Error rate
- Consistency across different reviewers and review teams

Resource Metrics

- Reduction in contractor spend or external review costs
- Efficient use of staff time and labor hours

- Ability to reallocate resources to higher-value work
- Sustainable spend across budget cycles

Strategic Metrics

- Enhanced accountability and greater confidence in meeting mission objectives
- Faster response times
- Increased staff satisfaction and retention
- Improved turnaround for high-volume requests and reviews

When defined and measured together, these metrics show the broader organizational impact and provide evidence that can be communicated to leadership and stakeholders.

Next Steps for Public Sector Leaders

Adopting generative AI in government is not about chasing hype. It is about applying technology responsibly in records-intensive work so agencies can strengthen transparency, uphold oversight, and deliver on their missions more effectively. How do you approach the introduction of generative AI with confidence?

To streamline and keep track of the process, follow these steps:

1. Identify a use case and internal champion
2. Engage vendors that demonstrate transparency, defensibility, security, compliance, and mission alignment
3. Develop a pilot plan with well-defined success metrics
4. Educate staff, document workflows, and refine approaches based on outcomes
5. Measure, validate, and iterate
6. Scale thoughtfully, prioritizing high-impact and mission-critical areas

Generative AI is already delivering measurable impact across government tasks. Experience shows that agencies build trust in AI by starting small, moving deliberately, and scaling what demonstrates value.

Relativity aiR for Review is primed for government agencies. It helps reduce backlogs, meet statutory deadlines, and improve mission outcomes without increasing headcount or compromising oversight. Tools like Relativity aiR for Review, grounded in transparency, security, and accountability, prove that generative AI can become a trusted tool for government service – and the agencies that scale responsibly will set the standard for how it delivers lasting public value.



231 South LaSalle Street | 20th Floor | Chicago, Illinois 60604

+1 (312) 263-1177 | relativity.com

© Relativity. All rights reserved.