HALYARD CONSULTING

Designing Ethical AI Workflows for Government & Public Institutions.

Created by: Halyard Consulting
Date: August 5, 2025
Written by: Jonathan Goodman

Executive Summary

Artificial intelligence (AI) is rapidly transforming how governments, public agencies, and institutions operate, streamlining service delivery, increasing responsiveness, and enabling data-informed decision-making. But with great capability comes great responsibility.

For public-sector leaders, the challenge isn't just adopting AI, it's doing so ethically, transparently, and equitably. Unlike in the private sector, AI adoption in government must navigate complex public interests: civil liberties, accessibility, bias mitigation, due process, and community trust.
\
This whitepaper provides a roadmap for designing **ethical AI workflows** tailored to the unique needs of public institutions. We explore how government agencies can responsibly implement AI without compromising public trust or compliance. Drawing on our work with municipalities, education systems, and infrastructure partners, we offer strategies that balance **innovation with governance, automation with accountability,** and **efficiency with equity.**

You'll learn how to:

- Align AI workflows with public values and legal mandates.
- Identify ethical risks across the AI lifecycle.
- Implement safeguards such as human-in-the-loop review and auditability.
- Design inclusive, multilingual, and ADA-compliant systems.
- Use AI not just to streamline government, but to strengthen it.

At Halyard Consulting, public institutions deserve AI tools built not just for outcomes, but for fairness. This whitepaper helps decision-makers lead with both confidence and conscience.

The Public Sector AI Landscape & Imperative for Ethics

Artificial intelligence is no longer a speculative technology in the public sector; it's already being used to manage traffic systems, triage benefits claims, translate public notices, and automate constituent engagement. As agencies face increasing pressure to "do more with less," AI presents a powerful lever for modernization and efficiency.

But in contrast to the private sector, where profit and scale often drive adoption, government entities must prioritize **fairness, transparency, and the public good.**

🚨 Why Ethics in Government AI Matters

- **Public Accountability:** Government systems must uphold constitutional and legal protections, including due process, civil rights, and accessibility.
- **Power Imbalance:** Unlike private services, citizens often cannot "opt out" of government systems. This raises the stakes for fairness and accuracy.
- **Bias & Discrimination Risks:** Poorly designed AI can reproduce or even amplify historical inequities, particularly in housing, public safety, and social services.
- **Trust in Institutions:** Public skepticism toward automation is higher when applied to enforcement, eligibility, or decision-making. Transparent workflows are essential.

📈 Trends Driving AI Adoption in Government

- **Federal & State Modernization Mandates:** Executive orders (e.g., EO 13960 & EO 14110) now require responsible AI practices for agencies receiving federal funds.
- **Labor Shortages & Backlogs:** AI agents and workflow automation are easing administrative burdens and helping teams serve more constituents with fewer resources.
- **Multilingual & Inclusive Communication Needs:** AI-powered translation and accessibility tools are bridging gaps for non-English speakers and people with disabilities.
- **Data-Driven Policy Goals:** From Vision Zero traffic programs to climate resilience dashboards, AI enables actionable insight across departments.

Government leaders are not asking whether to use AI, but **how to do it responsibly.** The following sections of this whitepaper walk through the building blocks of ethical design, covering everything from values-based workflow planning to implementation safeguards.

Principles of Ethical AI in Government Workflows

Designing ethical AI systems in the public sector isn't just a technical process; it's a civic responsibility. AI must operate within a framework that reflects democratic values, legal mandates, and community needs.
The following principles provide a foundation for building trust, ensuring compliance, and preventing harm.

1. Transparency

"If a system makes a decision, the public deserves to understand how and why."

- Use explainable models when feasible, especially when outcomes affect benefits, enforcement, or eligibility.
- Publish documentation of training data sources, model purpose, and decision criteria.
- Label automated interactions clearly (e.g., "This message was generated by an AI assistant").

2. Accountability

"AI decisions should never be final without human oversight."

- Establish a chain of responsibility for AI outputs and actions.
- Implement a human-in-the-loop review where outcomes affect legal or financial standing.
- Log decision paths and provide appeal processes for impacted constituents.

3. Fairness & Equity

"AI must work for everyone, especially the historically underserved."

- Audit for disparate impact across race, gender, language, disability, and income.
- Ensure datasets reflect the communities the system will serve.
- Co-design workflows with marginalized stakeholders, not just internal staff.

Principles of Ethical AI in Government Workflows (cont)

4. Privacy & Data Stewardship

"The government holds sensitive data in the public trust; it must be handled as such."

- Minimize data collection to only what's necessary for the task.
- Anonymize personally identifiable information (PII) wherever possible.
- Align with state and federal privacy laws (e.g., HIPAA, FERPA, CJIS, CPRA).

5. Accessibility & Inclusion

"If a resident can't use it, it doesn't serve the public."

- Ensure AI interfaces are WCAG-compliant, multilingual, and mobile-friendly.
- Avoid tools that assume digital literacy or broadband access.
- Offer parallel non-AI options when needed (e.g., call centers, printed forms).

These principles aren't aspirational—they're actionable. The following section will explore **how to operationalize these values across the AI implementation lifecycle**, from project scoping to post-deployment monitoring.

Designing AI Workflows with Ethics by Design

Ethics can't be an afterthought; it must be embedded in every step of your AI project lifecycle. From scoping and procurement to deployment and maintenance, each stage is an opportunity to build transparency, equity, and public trust into your systems.

At Halyard Consulting, we use a methodology called **Ethics by Design**, which maps values directly into workflows and decision points.

🧭 Step 1: Define Purpose & Public Impact

- What problem is the AI meant to solve?
- Who benefits and who could be burdened?
- What public values (e.g., equity, efficiency, accessibility) does this support?

📌 Example: A city chatbot isn't just answering FAQs, it's shaping the public's perception of municipal responsiveness.

📝 Step 2: Select Use Cases Carefully

Not all government tasks are appropriate for AI automation. Good candidates share these characteristics:

- Repetitive, rules-based logic (e.g., appointment booking, permit lookup).
- Low risk of rights infringement.
- High administrative burden with low current responsiveness.

🚫 Avoid: Eligibility determination, law enforcement, or disciplinary workflows without strong guardrails.

👫 Step 3: Include Diverse Stakeholders Early

- Invite residents, advocacy groups, and front-line staff into co-design sessions.
- Pay special attention to populations most likely to be excluded or misrepresented by technology.
- Document feedback and update use cases accordingly.

Designing AI Workflows with Ethics by Design (cont)

📐 Step 4: Design for Explainability

- Choose tools that allow outputs to be logged, audited, and explained.
- Use decision trees, prompt libraries, and fixed logic where black-box models pose risks.
- Include fallback paths that escalate uncertain cases to a human agent.

🔍 Step 5: Test for Bias & Functionality

- Simulate edge cases across demographics, device types, and languages.
- Test with screen readers and older browsers.
- Run small-scale pilots before city-wide rollouts.

🔄 Step 6: Monitor, Maintain & Adapt

- Set up real-time usage analytics and public feedback loops.
- Appoint internal “AI stewards” to oversee updates and performance reviews.
- Deactivate or modify workflows that show harmful or unequal patterns.

Building an ethical AI workflow doesn’t require perfection; it requires intention, participation, and continuous refinement.

Ethical AI in Action: Real-World Government Use Cases

The following anonymized examples illustrate how public-sector agencies have deployed AI workflows while prioritizing transparency, equity, and trust. Each showcases how ethical design can be both practical and powerful when applied intentionally.

🛣️ 1. Vision Zero Engagement Assistant — City Transportation Department

Challenge:
A midsize city wanted to collect resident feedback on traffic safety improvements in historically underserved neighborhoods but lacked the capacity to manage multilingual public input across channels.

Solution:
A multilingual AI chatbot was deployed on the city's Vision Zero website and integrated with SMS outreach. It allowed residents to submit feedback on unsafe intersections and request pedestrian improvements in English, Spanish, and Tagalog.

Ethical Features:

- Residents informed when they were interacting with AI.
- All submissions exported weekly to a human analyst for prioritization.
- Community groups helped shape the bot's tone and FAQs.

Impact:

- 3× increase in neighborhood-specific traffic concerns submitted.
- Higher engagement from non-English-speaking households.
- Faster data synthesis for traffic engineering teams.

Ethical AI in Action: Real-World Government Use Cases (cont)

🏛️ 2. Court Reminder Voice Assistant — County Judicial System

Challenge:
A rural county court experienced high failure-to-appear (FTA) rates due to missed hearing notifications and language barriers.

Solution:
An outbound voice assistant was deployed to remind residents of upcoming court dates, verify attendance, and connect them to legal aid. Messages were delivered in the recipient's preferred language based on intake records.

Ethical Features:

- No AI was used for legal advice or eligibility decisions.
- Opt-out option clearly communicated.
- All voice scripts were reviewed by public defenders and accessibility experts.

Impact:

- 22% reduction in FTAs within 90 days.
- Increased trust among residents who had previously disengaged with court communication.
- Cost savings from fewer rescheduled hearings.

Ethical AI in Action: Real-World Government Use Cases (cont)

📚 3. School District Chatbot — Urban K-12 System

Challenge:
A large urban district struggled with fielding thousands of repetitive parent inquiries during enrollment season, leading to long wait times and frustration.

Solution:
A plain-language chatbot was introduced on the enrollment portal, answering FAQs about school choice, transportation, documentation, and deadlines.

Ethical Features:

- Designed with accessibility consultants to support screen readers and mobile-first usage.
-
- Offered translation in 5 languages.
- Escalated unclear queries to human staff during business hours.

Impact:

- Reduced email and phone volume by 40%.
- Improved family satisfaction with enrollment communication.
- Equitable access for non-native English speakers.

These examples demonstrate that ethics and automation are not opposites, they're **mutually reinforcing**. When designed with community input and oversight, AI can strengthen public services and civic trust.

Checklist for Ethical AI Procurement & Implementation

Whether issuing an RFP or evaluating a vendor proposal, public institutions must hold AI tools to a higher standard. This checklist provides a practical framework for evaluating whether a system aligns with your agency's ethical, operational, and legal obligations.

✅ **Use this checklist during planning, procurement, or internal audits.**

🧩 Strategy & Use Case Fit

- Is the AI being applied to a task that is repetitive, rules-based, and well-bounded?
- Have alternative non-AI approaches been considered?
- Does the use case advance a public value (e.g., equity, efficiency, access)?

👥 Stakeholder Engagement

- Have impacted communities or civil society groups been consulted?
- Were frontline staff involved in defining success criteria and edge cases?
- Is there a plan to collect ongoing resident feedback?

🔍 Transparency & Explainability

- Are AI-driven decisions explainable in plain language?
- Are residents clearly notified when they're interacting with AI?
- Can outputs be audited or traced back to decision logic?

⚖️ Fairness & Bias Mitigation

- Has the tool been tested across diverse demographic groups?
- Were training datasets evaluated for representativeness and bias?
- Are there safeguards against disproportionate impact?

Checklist for Ethical AI Procurement & Implementation (Cont)

🔐 Privacy & Security

- Is data collection limited to what is strictly necessary?
- Is all personally identifiable information (PII) encrypted or anonymized?
- Are data governance practices documented and compliant with applicable laws?

♿ Accessibility & Inclusion

- Is the interface WCAG 2.1-compliant?
- Does the tool function on mobile, screen readers, and older browsers?
- Are multilingual and plain-language options provided?

🔁 Oversight & Governance

- Is there a human-in-the-loop or escalation path for high-impact decisions?
- Are logs and audit trails retained?
- Has a staff member or team been assigned as the AI system owner?

Checking these boxes isn't just risk management, it's public service. Agencies that prioritize ethical workflows are better positioned to serve diverse communities, build long-term trust, and avoid costly rollbacks or litigation.

Conclusion: Leading with Ethics in the Age of Automation

Artificial intelligence offers extraordinary potential for government transformation, but only if it is grounded in public values. The greatest risks with AI in the public sector aren't technical, they're ethical.

The good news? Ethical AI isn't an obstacle to innovation. It's the foundation of meaningful, sustainable progress. When designed responsibly, AI doesn't just streamline public services, it makes them smarter, fairer, and more responsive to the people they serve.

From transportation planning to court reminders, community engagement to internal automation, the opportunity is clear: agencies can deliver faster, more accessible, and more inclusive services, without sacrificing accountability or public trust.

What's required is leadership. Not just from technologists, but from program managers, policy directors, and civic decision-makers who understand that **how** we implement AI is just as important as **what** it does.

This whitepaper provides a roadmap for designing those workflows. And as more institutions take this step, ethical AI can become the norm, not the exception, in public sector innovation.

About Halyard Consulting

Halyard Consulting is a strategy-first consulting firm that helps public agencies, school systems, and mission-driven organizations modernize operations and improve outcomes through responsible AI. Our work blends deep consulting experience with implementation expertise to deliver real-world solutions that are accessible, transparent, and equitable.
We specialize in:

- AI policy and ethical governance frameworks.
- Voicebot and chatbot deployment for multilingual public services.
- Community engagement automation and accessibility tools.
- AI-powered workflow automation for government operations.

As a certified LGBTQ+, Disability-Owned (DOBE), and New Jersey Small Business Enterprise (SBE), Halyard Consulting brings a unique perspective to civic technology, rooted in inclusion, accountability, and public value.

Let's Build Better Together

If you're a public-sector leader looking to responsibly implement AI, we'd love to talk.
📅 Schedule a Consultation
🌐 Explore Our Tools
📥 **Download More Resources**