PayIt™

Getting AI-Ready:

A Practical Guide for State and Local Government Leaders

Contents

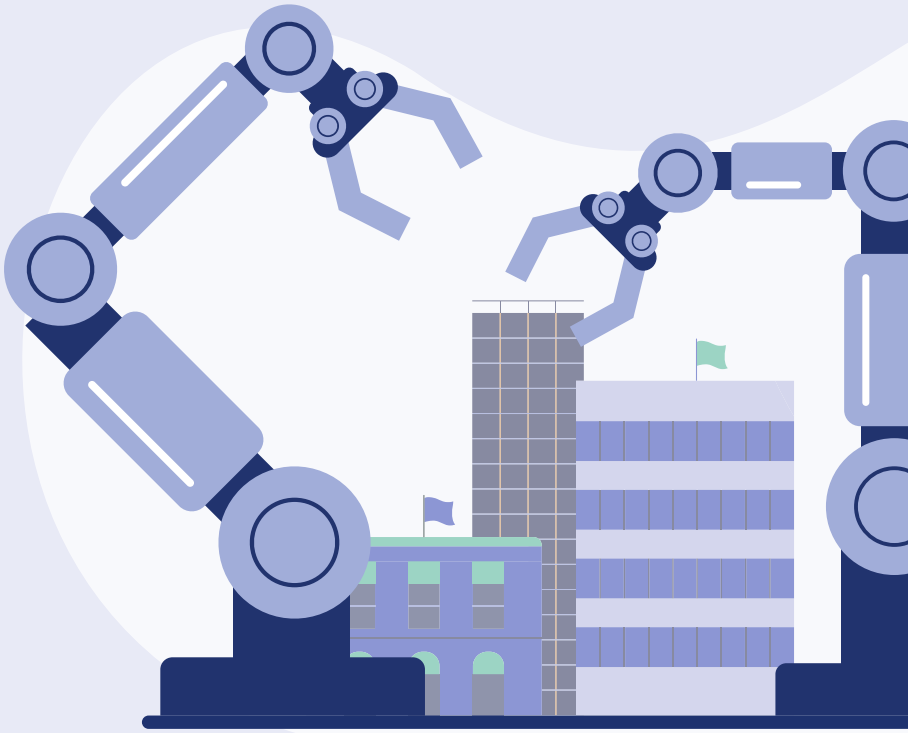

AI is moving from intention to implementation

Government leaders face real pressure to act on AI — and the agencies that get the most out of it will be the ones that spend time on the right preconditions before they deploy.

According to PayIt's 2026 Government Payments Experience Index, 64% of state and local government leaders believe AI will positively impact their ability to serve residents faster, and 63% see it as a path to limiting tax increases through improved operational efficiency. Those attitudes are driving concrete activity: 34% of agencies have already deployed AI for back-office operations and 26% for resident-facing services, while another 50% are actively evaluating back-office AI and 6% are currently in implementation. Resident-facing activity is running at a similar pace, with 58% evaluating and 6% in implementation. Across back-office and resident-facing functions alike, AI is firmly on the agenda for state and local government.

State and local governments tend to take a measured approach to adopting new technologies — particularly those with significant implications for data security, staffing, and public accountability. That deliberate pace reflects the stakes involved, not a lack of urgency.

AI and data analytics have ranked consistently among the top priorities in NASCIO's annual State CIO Survey — a reflection of how broadly state technology leaders see AI reshaping government operations. And among agencies that have already implemented AI solutions, 93% report a positive or very positive impact on operating efficiency.

This guide is written for agency leaders who have AI on their agenda — who recognize the potential and want to get the preconditions right before committing to a deployment. It covers the six areas that most often determine whether an AI initiative delivers: use case selection, data readiness, , change management, governance, resident communication, and cost planning.

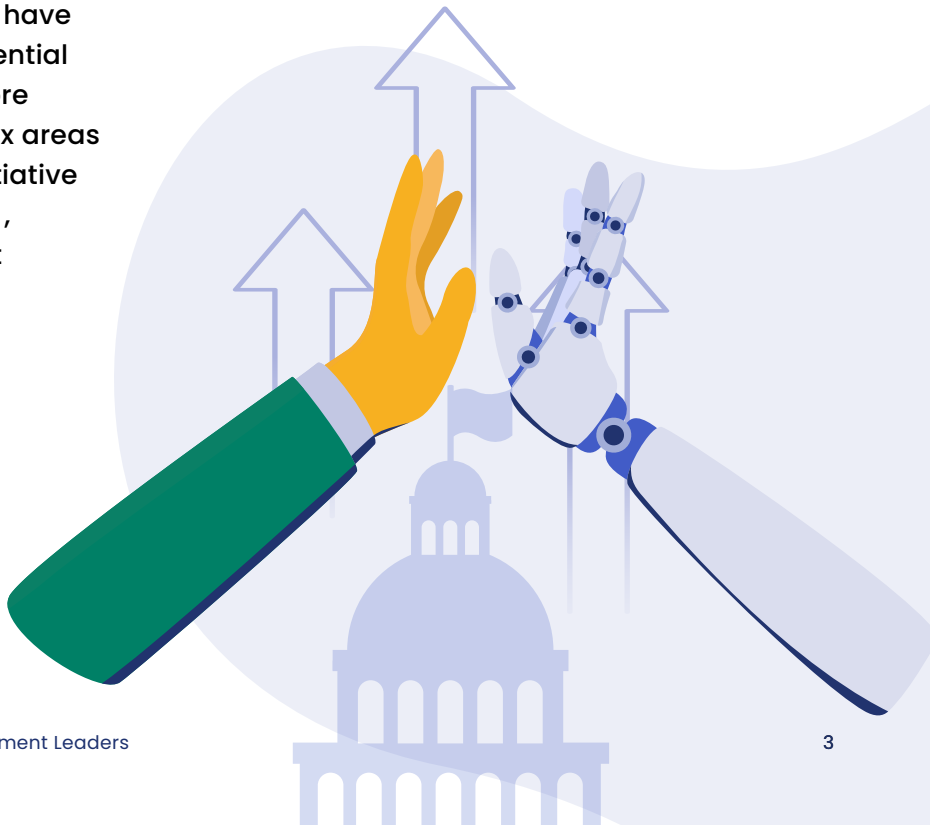

What AI can (and can't) do for your agency

Agencies that start by identifying their most pressing operational problems — those with enough scale to be worth automating and enough structured data to support AI — consistently get better results than those that start with a technology and look for problems to solve.

AI performs consistently and at scale on specific categories of work; processing high volumes of structured tasks, translating natural language into queries or actions, verifying documents and identities, and helping people find and act on complex information. How well it performs depends heavily on data quality: clean, structured data yields reliable results, while fragmented or inconsistent data compounds over time into errors.

Consequential decisions about individual residents — eligibility determinations, penalty assessments, appeals — require human oversight as a matter of policy, not just good practice. Government decisions carry due process obligations, and those obligations need a person in the accountability chain. Build that into your workflow design from the start.

For state and local government agencies, the most practical near-term applications fall into four areas:

Resident-facing information and interaction.
Residents increasingly approach government services the same way they approach everything else: with natural language questions typed or spoken in full sentences. "What do I need to renew my vehicle registration?" "My property tax assessment went up — how do I appeal?" AI-powered chat and answer tools handle these queries at any volume, at any hour, and substantially reduce call center and counter traffic. The value compounds as residents adopt the digital channel: fewer paper submissions, fewer in-person transactions, and simpler reconciliation for staff.

Staff-facing information and navigation.
Government staff face their own version of the information access problem: finding the right answer often means navigating multiple databases, years of policy updates, and regulatory requirements that may not be well-documented in any single place. AI tools built on an agency's own knowledge base — procedures, regulations, policy histories, and system documentation — can surface answers to staff questions quickly and consistently, reducing time spent searching across systems and the risk of working from outdated information. The same capability that helps a resident understand a billing question can help a caseworker identify the right exemption criteria or locate the correct procedure for an unusual filing.

Document and identity verification.
Government agencies process enormous volumes of documents: permit applications, vital records requests, license renewals, court filings, and benefit submissions (to name a few). Each one requires verification: is the document complete? Are the required fields legible and consistent? For identity verification specifically, AI can match a submitted selfie to the driver's license photo on file, automating a check that would otherwise require manual review on every submission. Done manually at scale, this work is a significant drag on processing times and staff capacity. AI handles verification accurately and quickly, flagging exceptions for human review while automatically clearing the straightforward cases.

Back-office process automation.
Data entry from fragmented payment and case management systems, application status tracking across departments, financial reconciliation across funding streams, and revenue forecasting are all rule-based processes well-suited for AI assistance. These tasks are time-consuming and error-prone when done manually, and rarely require the kind of judgment that staff should be spending their time on. Automating them frees capacity for higher-value work.

Choosing your first use cases

Keep in mind that agencies shifting residents to digital channels don't just improve the resident experience — they also reduce paper processing, cut phone calls and counter traffic, and simplify reconciliation. The efficiency gains compound as digital adoption grows. AI-powered resident tools are among the most effective ways to drive that shift, making use case selection both an operational and a strategic decision.

Not all use cases are equally ready or equally valuable. A first deployment that struggles creates organizational skepticism that's hard to undo. A first deployment that works builds confidence and justifies budget support to go further. Evaluate potential use cases against four criteria:

Staff time and resources recovered.
Does this task consume significant staff time or budget? Strong early candidates are processes with predictable, rule-based steps — document verification, common inquiry response, status tracking — where AI can handle the volume and flag exceptions for human review. Human oversight isn't a binary condition: a deployment where AI handles 70% of cases and routes the remainder to staff still recovers meaningful capacity, and that partial automation often makes sense for cases that involve some degree of judgment. Cases that require sustained relationship management or legal accountability may not be fully automatable yet, but that's rarely a reason to do nothing.

Data availability.
Is the data this use case needs already clean and accessible? If yes, you can move quickly. If not, data preparation comes first. This criterion often determines the sequencing of your deployment roadmap more than anything else. Depending on the use case, your vendor may be able to help with data access and structuring.

Resident impact.
Does this use case meaningfully improve the experience of completing a payment, submitting an application, or getting an answer to a question? The strongest early use cases have both operational and resident-facing value.

Risk if the AI is wrong.
What happens when the system makes a mistake? A chatbot that gives an incomplete answer about a filing deadline is a recoverable error. An automated determination affecting someone's payment plan or permit status is not. Start where the stakes for error are lower, and build human review into workflows where they are higher.

For property tax agencies, AI-powered answer tools that help residents understand their assessment, find payment options, or navigate an appeal are natural starting points. The data is well-structured, the questions are predictable, and the volume — especially around billing cycles — is high enough to make the efficiency gain significant.

For licensing and permitting, document verification is a strong early use case. Checking whether an application is complete, documents are legible and consistent, and identity verification is valid substantially reduces manual review burden and speeds processing times for applicants when they're most likely to be frustrated by delays.

Getting your data house in order

State and local agencies rarely face a data shortage. What they typically lack is data that's connected, consistent, and accessible in a way that AI applications can use directly. Payment records sit in a financial system. Application status lives in a case management platform. Property records are in a land management system. Identity verification may be handled separately from any of those. Each system was chosen for a specific operational purpose and is the authoritative source for its slice of information.

That's the typical starting point for AI readiness — not one tidy database, but many purpose-built ones that don't share a common structure or communicate easily with each other. Data fragmentation is not a government-specific problem; large enterprises across every industry are navigating the same challenge. But it does mean that the path to AI readiness runs through data integration, not just tool selection.

Getting ready for AI means moving through three stages: an assessment of what you have, an action plan to address gaps, and connecting your data to the AI tools you plan to deploy. Skipping the first two stages leads to stalled or unsuccessful deployments.

A data readiness assessment for each priority use case should answer four questions:

1. **Where does the relevant data live, and which system is the authoritative source for each data type?**
   A property tax AI might need assessment data from a tax system, payment history from a financial platform, ownership records from land records, and exemption status from a benefits system — each potentially in a different place.

2. **How consistent are the data structures?**
   Data formats vary more than most teams expect; dates formatted as MM/DD/YYYY in one system and YYYY-MM-DD in another, street addresses that abbreviate in some records and spell out in others, name fields with varying capitalization that cause mismatches in automated lookups. A simple spreadsheet export and comparison across systems can surface these patterns quickly. Inconsistency at scale defeats AI accuracy.

3. **How complete is the data?**
   Pull a random sample of 10–15% of your records and check each one for the specific fields your AI application requires. If a required field is blank or unusable in a substantial share of sampled records, that gap needs to be addressed before deployment. The sample gives you a reliable proxy for what you'll find across the full dataset.

4. **How current is the data?**
   Outdated data produces wrong answers. Talk to the teams that update each system and understand whether changes flow in automatically, manually, or inconsistently.

The practical path for most agencies isn't solving the entire data infrastructure problem before deploying AI. It's a use-case-by-use-case approach: for each planned AI deployment, identify the specific data fields it needs, trace those fields to their source systems, and build or obtain the extraction capability for those fields specifically. This keeps early deployments moving without requiring a multi-year data transformation project as a prerequisite.

Legacy systems create a specific complication: they often lack modern interfaces, can't export data in standard formats, and require custom extraction work. If your source system for a priority use case is a legacy platform, that extraction work is usually the first project on the list — not the AI deployment itself. Vendors with experience in government technology can often accelerate this, having built connectors for common legacy government platforms. It's worth asking about this during procurement.

Some agencies are getting ahead of the fragmentation challenge by building centralized data environments — often called data lakes or data platforms — that consolidate records from multiple source systems into a single queryable layer. This doesn't replace the source systems; it creates a synchronized copy that AI applications and analytics tools can access without disrupting operational workflows. Cloud-based platforms have made this more accessible than it was five years ago, and several government technology vendors offer pre-built connectors for common government systems.

For agencies with well-structured databases and privacy-sensitive data, an emerging architectural approach worth noting is AI that never touches your data directly. Crawford County, Pennsylvania, demonstrated this at the 2026 NASCIO Midyear Conference. Their parcel viewer AI tool translates natural language questions ("find properties within a half-mile that sold for more than $1 million in the last three years") into database queries and returns results without storing or processing the underlying resident data. For some use cases, this architecture addresses privacy concerns while simultaneously reducing data preparation burden.

Build vs. buy: the question is back on the table

The past decade of modernization has given state and local agencies important tools: cloud-based platforms that can be configured to government needs without the cost and burden of building and maintaining custom software. Most agencies moved in that direction for good reason.

Off-the-shelf AI tools, including readily available APIs from major providers, have reopened the conversation about building solutions in-house. Some IT teams see an opportunity to develop custom applications with more control and flexibility. The rise of vibe coding tools in particular may prompt the question: "can't we just build this ourselves with AI?"

But the hidden costs of self-building tend to accumulate faster than initial estimates suggest. A custom AI application built on a general-purpose model won't come pre-configured for government data structures, compliance requirements, or the specific transaction workflows your agency runs. You have to build that institutional knowledge from scratch, maintain it as models evolve, and own all the security and accuracy obligations that come with it. And even as AI makes it possible for non-technical people to build applications, the sensitivity of most government data and processes demands stringent security protections and safeguards that vibe coding tools do not provide.

Purpose-built platforms designed for government organizations offer a different path. For agencies adopting a still-nascent and genuinely consequential technology, the advantages matter:

- Faster time to value: No multi-year build before you see results
- Security and compliance managed by the vendor: Certifications, patches, and data handling obligations are maintained continuously
- Continuous innovation: Model improvements and new capabilities ship to all customers without requiring internal IT projects
- Government context built in: Workflows, edge cases, and compliance requirements specific to property tax, licensing, courts, and utilities are already accounted for
- Contractual accountability: When something goes wrong, there's a defined remediation path

For government agencies in particular, working with a trusted vendor means you don't have to become an AI company to benefit from AI. The right question isn't whether your team ***could*** build it. It's whether building and maintaining it for the next five years is the best use of your capacity.

Getting your team ready

The people using your AI technology are what determine whether a deployment delivers on its promise. Tools that staff don't trust, don't understand, or don't know how to use correctly produce worse outcomes than the manual processes they replaced. Getting the team dimension right is as important as getting the technology right.

Government workforces have specific concerns about AI that warrant direct and early attention — job security is the most common. Some tasks will change significantly, and government staff are navigating the same shift that knowledge workers across every industry are working through: roles are evolving to include directing, reviewing, and interpreting AI-generated work alongside doing the work directly. Agencies that engage with that reality early, and invest in helping staff build those new skills, see faster adoption and fewer implementation problems. That conversation is more effective before deployment than after resistance has already set in.

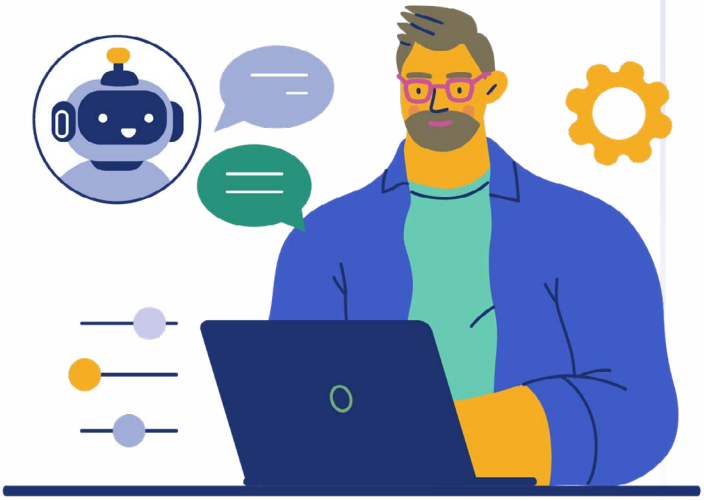

Six things that help prepare your team for an AI shift:

- **Deploy AI as policy with human review.**
  Staff are more willing to engage with AI tools when they know it's there to assist them, they retain authority over final decisions, and the rules of engagement are governed by clear policy. For any process that affects resident accounts, eligibility, penalties, or legal matters, document the human review step in the workflow before launch.
- **Be direct about what will change**.
  If a tool eliminates a task someone currently does, name it plainly and explain what that person will do instead. Vagueness creates more anxiety than the change itself.
- **Build an AI task force from your early adopters.**
  Most agencies have people who are curious and willing to try new things. Give them early access, formalize them as a small cross-functional group, and charge them with piloting tools, documenting what works, and developing training for their peers. Peer-led training is more effective than vendor-led training for skeptical staff — colleagues doing the same work have more credibility than an outside trainer.
- **Train for the actual workflow.**
  Generic AI literacy training isn't enough. Staff need to understand specifically what the tool does in their day-to-day context, what its failure modes look like, and exactly how to handle exceptions. Build scenario-based training around the real work, not abstractions.
- **Create a feedback loop in the first 90 days.**
  Build a simple mechanism for staff to report what's working and what isn't — a short weekly check-in, a shared log, anything that captures signal early. Problems caught in the first 90 days are easier to fix and signal to staff that leadership is paying attention.
- **Emphasize leadership visibility.**
  Directors and department heads who visibly use and reference AI tools signal that adoption is real and expected — not a directive from IT that stops at the department door.

Bringing residents along

Change management doesn't stop at the staff level. Residents interact with the output of your AI systems, and how they feel about that interaction affects both their trust in your agency and their willingness to use your digital channels.

PayIt's 2025 Consumer Digital Government Adoption Index found that 43% of residents are more concerned about state and local government's use of AI than the private sector when it comes to keeping personal information safe. Nearly 4 in 10 (39%) are more concerned about government AI providing incorrect or incomplete information.

Those concerns reflect something real: haphazard AI use in government carries consequences that a bad product recommendation doesn't. Residents who feel uncertain or blindsided by AI are less likely to complete digital transactions and more likely to call or come in — the opposite of the efficiency outcome you're working toward. Agencies that get ahead of resident concerns, rather than waiting for them to surface, earn trust faster.

Resident comfort varies significantly with the type of AI interaction. People are most at ease with AI that helps them find and act on information: 69% are at least somewhat comfortable with AI predicting and surfacing relevant information during a government website visit, and 65% are comfortable with AI answering questions about services, policies, or rules. Comfort drops considerably when AI plays a decision-making role: only 55% are at least somewhat comfortable with AI evaluating their eligibility for a government benefit. (2025 Consumer Digital Government Adoption Index)

Starting with AI that informs and assists rather than AI that makes decisions is not only good policy, it's a good communication strategy — it builds resident familiarity with AI in government before you ask them to accept its role in higher-stakes processes.

Here are a few steps to keep resident confidence high:

Disclose when residents are interacting with AI. A simple acknowledgment — "This chat is powered by AI. A staff member can help with more detailed requests" — is reassuring, not alarming. Disclosure builds credibility; the absence of it erodes it.

Explain what safeguards are in place. Residents want to know their data is protected, that humans review consequential decisions, and that there's a path to reach a person when the AI can't help. Make those things visible in the interface and easy to find on your website.

Build a feedback mechanism into the resident experience. A simple post-interaction rating, a feedback button on chatbot responses, or a regular review of support tickets related to AI interactions gives you valuable information about where residents are struggling. Agencies that listen and respond to that feedback — and say so publicly — demonstrate accountability in a way that static disclosures don't.

Communicate proactively at scale. Include a plain-language explanation of how your agency uses AI in your website FAQ, in any digital communications related to AI-assisted processes, and in your general resident communications cadence. Don't wait for a complaint or a news story to prompt the conversation.

Governance, policy, and accountability

The governance conversation needs to happen before deployment. Agencies that deploy first and figure out policy afterward tend to find themselves in harder conversations with auditors, elected officials, or residents than if the guardrails had been designed in from the start. The good news is that frameworks and precedents are accumulating quickly.

The NIST AI Risk Management Framework is a widely referenced starting point for government AI governance. It provides a structured approach to identifying, assessing, and managing AI risk across four functions: Govern, Map, Measure, and Manage. It's voluntary, but federal agencies are beginning to use it as a baseline, and many states and localities have adopted it as a reference for their own policies.

At the state level, several have enacted AI governance legislation requiring agencies to maintain inventories of AI systems, conduct impact assessments on high-risk deployments, and publicly disclose how AI is being used. The Center for Democracy and Technology and MRSC maintain current resources on local government

Your governance framework should address five core questions:

Where is AI being used, and for what purpose?
Maintain an inventory of AI applications, the data they use, and the decisions they influence. This is the foundation for every other accountability mechanism.

What requires human review before action is taken?
Define clearly which AI-assisted processes need staff sign-off before an output is deployed. This is especially important for anything affecting resident accounts, penalties, eligibility, or legal matters — and it should be documented formally, not left to individual judgment.

Can you explain AI-influenced decisions to residents?
Government agencies carry due process and transparency obligations that private-sector organizations don't. When AI influences a determination that affects residents, they should have a right to know AI was involved, understand how the determination was made, and have avenues to contest it. Before you deploy, ensure you have a clear disclosure of AI's role in resident-facing communications in place, and have a documented pathway for residents to challenge determinations made with the support of AI.

How are errors identified and corrected?
AI systems make mistakes. Define the process for a staff member or resident to dispute or flag an error, have it reviewed by your staff, and get it corrected. Without a defined review path — particularly for residents impacted by AI mistakes — trust erodes and accountability gaps grow.

What security and privacy protections are in place?
Any AI-powered solution should meet the highest standards for security and the protection of private data including personally identifiable information (PII) and payment details. Ensure that multiple layers of security are in place, and that your team and vendors have an active security posture. Understand how the underlying models you and your vendors use are trained, and what, if any, of your agency's data is used in that process.

It's important to keep in mind that if an AI-powered solution is delivered through a vendor, your contract should address data handling, model accuracy standards, audit access, and what happens when something goes wrong. These are negotiable terms — treat them that way.

Planning for the real costs of AI

AI pricing can work differently from traditional government software licensing, and agencies that don't plan for it may be surprised. Most AI applications are priced based on usage — the number of requests processed, the volume of data handled, and how often models are accessed — rather than a fixed annual seat license. Costs scale with volume, which means they look very different at 5,000 monthly resident interactions than at 50,000.

If you are purchasing an AI solution from a vendor, make sure you understand how it's priced. If the pricing is usage-based, ensure that there are safeguards in place to prevent unexpected usage bills. Finally, have a clear idea of how the vendor manages the drivers of AI costs outlined below:.

Inference costs. Every AI model request, from resident queries in a chat box to document verifications or identity checks, measures usage on a per-item cost (typically called 'tokens'). AI usage accumulates at scale and can spike quickly during high-volume periods, such as property tax billing cycles.

Data storage and processing. AI applications that work with documents, images, or large record sets need storage and processing infrastructure. Storage is generally predictable; processing costs depend on volume and frequency.

Integration and maintenance. Connecting AI tools to your existing systems requires up-front integration work and ongoing maintenance as those systems change. Budget for both.

Model updates and versioning. AI models evolve, and vendors release updates. Understand whether updates are automatic, optional, or require re-training — and what those updates cost.

Several strategies help contain costs without sacrificing effectiveness, whether you are using a vendor solution or your own:

Several strategies help contain costs without sacrificing effectiveness, whether you are using a vendor solution or your own:

- **Right-size the model to the task.** Powerful (and expensive) AI models aren't always necessary. For structured, rule-based tasks such as status lookups or document completeness checks, simpler and less costly approaches often perform just as well as more complex AI models.
- **Batch non-time-sensitive processing.** Requests that don't need real-time responses (such as overnight reconciliation runs or document reviews) can be processed in off-peak windows at lower cost.
- **Cache frequent queries.** If a large percentage of resident questions fall into predictable categories, a well-designed system can serve cached responses to common queries rather than running a full AI inference for each one.
- **Consider interface-layer AI for appropriate use cases.** For agencies with well-structured databases, AI that translates natural language into database queries — rather than ingesting and processing data directly — can reduce both inference costs and data privacy risk. (See the Crawford County example.)
- **Phase your rollout.** Starting with a single, well-defined use case lets you validate cost projections before scaling. Ask vendors explicitly what happens to costs if usage doubles — that answer tells you whether the model works at your expected scale.

The efficiency gains from well-chosen AI deployments typically outpace the costs. That case is strongest when the costs are modeled accurately from the start and built into your budget request alongside the projected savings.

Starting your readiness assessment

AI readiness is a set of conditions, some you have today and some you're building toward. Every organization, from government agencies to large enterprises, is managing through readiness gaps. A useful internal review covers five dimensions:

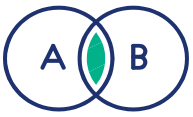

Use cases. Have you identified two to three specific, high-efficiency problems to address first, evaluated against the criteria above, and confirmed that the data is properly formatted and accessible?

Data. Do you know where the data for your priority use cases lives, which systems house it, and what its quality looks like? Have you identified the extraction path for each?

Team readiness. Do you have a plan for preparing staff? Have you identified a pool of early adopters, identified workflow-specific trainings, and prepared a feedback mechanism for early adopters?

Resident communication. Do you have a plan to disclose AI use, explain safeguards, and establish channels for resident feedback?

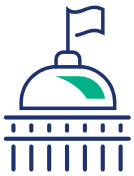

Governance. Do you have a policy framework in place — or a concrete plan to build one — that covers your AI inventory, human review requirements, error correction, and vendor obligations?

Think about each of the five readiness conditions. If you check most of the boxes, you're likely ready to start piloting AI in your agency. Choose one of your use cases, define what success looks like, and start to put it in place. If you're a mix of prepared and not quite ready, that's useful information. Identify the gaps in your readiness, and develop a framework to overcome them, from a small, defensible use case to preparing data for piloting with an early adopter team.

Among agencies already using AI, 72% plan to expand its usage into new areas over the next 12 to 18 months (2026 Government Payments Experience Index). The agencies that will make the most of this expansion are already laying the groundwork — but agencies who are just getting started can catch up as they take the steps outlined here.

AI readiness isn't a finish line. It's a working state while AI becomes part of how your agency operates and serves residents. And the success factors outlined here become an ongoing practice in high-performing organizations as AI becomes integral to their modernization efforts.

About PayIt

PayIt enables state and local government agencies to deliver a great resident payments experience that accelerates the shift to digital. Agencies choose PayIt to better achieve their mission through improved operational efficiency, customer support, and resident satisfaction. Our solutions span property tax, courts, utilities, DMV, outdoors, and more. PayIt provides a single resident profile across agencies and jurisdictions, integrates into back-office and adjacent systems, and our team helps clients drive adoption of digital channels. Serving more than 150 million residents in North America, we have received awards from Fast Company and StateScoop, and have been listed in the GovTech 100 for 10 years and counting.

PayIt™

Smarter for government.
Easier for everyone.

To learn more, visit www.payitgov.com