CLOUDFLARE

The AI Blueprint: Four Use Cases Defining the Future of the Public Sector

Realizing the benefits of AI requires an understanding of its risks

Artificial Intelligence in the public sector has officially shifted from speculative "what-ifs" to the urgent demand for scalable implementation. America's AI Action Plan, for instance, encourages rapid development and distribution of new AI technology across every industry – including government, defense, and higher education.

For public sector leaders, the challenge is no longer just understanding the technology, but identifying where it can deliver the most immediate and measurable value without compromising public trust. And that depends on four key use cases enable secure AI interactions and reinforce trust.

Use Case 1: Control internal use of GenAI

ChatGPT. Gemini. Claude. Generative AI tools are available, but realizing the potential can be difficult and slow. Internal policies may not be fully developed, even for products officially rolled out to the workforce. In addition, employees may not understand or trust GenAI technologies – while others might use unauthorized tools that increase risk.

Recent studies of internal GenAI usage reveal concerning patterns:

- **85%** of employees use AI tools before IT can vet them.[1]
- **93%** admit to putting internal data into AI without approval.[2]
- **63%** of breached organizations have no AI governance policies.[1]

Therefore, it's essential to control the internal use of GenAI. The first step is to identify and inventory the AI tools used throughout the organization – whether approved or not. Then limit access to sanctioned apps and ensure that the workforce uses them appropriately.

Modern technologies can help provide the necessary visibility and control:

- **Secure Web Gateway (SWG)** is a cloud security capability that connects workforce users to applications, providing essential visibility and control over GenAI activity in both sanctioned and unsanctioned tools.

- **Cloud Access Security Broker (CASB)** provides high-fidelity detail about SaaS environments, including sensitive data visibility and suspicious user activity. When integrated with an Identity provider, CASB offers insight into out-of-band, third-party AI usage.

Controlling internal use of GenAI is an essential use case to solve, but many public sector organizations are also planning to leverage AI externally to improve customer service – leading to the next use case.

Major AI use cases

1. Control internal use of Generative AI
2. Defend AI-powered apps from cyber threats
3. Protect interactions between AI agents and internal data
4. Build innovative AI apps and digital assistants

"Most agencies have started with internal, employee-facing chatbots to ensure they're comfortable with the technology and happy with the outcomes before creating external, customer-facing AI applications."

Dan Kent
Field CTO for Public Sector, Cloudflare

[1] 2025 ManageEngine research
[2] 2025 IBM Cost of a Data Breach report

Use Case 2: Defend AI-powered apps from cyber threats

What happens when a prompt injection attack on a car dealership's AI chatbot nets the attacker a $50,000 vehicle for just $1?

In the public sector, financial risks are just the beginning. AI exploits on public-facing AI applications can disrupt critical services, leak sensitive information, or cause the model to return incorrect, illegal, or harmful information. Therefore, a "safety-first" AI architecture is paramount for every public sector organization.

The four key recommendations for defending AI-powered apps from cyber threats are:

- **Inventory AI components**. Discover and track all AI components across web properties – from AI models to the APIs that integrate into backend systems – and keep the inventory up to date.
- **Protect AI models from abuse**. Evaluate and deploy modern AI firewalls designed specifically to detect and stop threats like prompt injection, model poisoning, and excessive usage that can bypass traditional security defenses.
- **Control sensitive data flow**. Analyze and control the bi-directional flow of sensitive information, both submitted by the user and returned by the AI model.
- **Deploy content guardrails**. Evaluate AI platforms that not only have integrated models, but also built-in guardrails that help prevent unsafe or toxic prompts from exploiting internal systems.

Modern AI Gateway solutions can help meet these recommendations. AI Gateways are proxy services that broker connections between the AI application and model providers (OpenAI, Anthropic, etc) with safety guardrails to ensure a consistent and safe experience, regardless of the model or provider used.

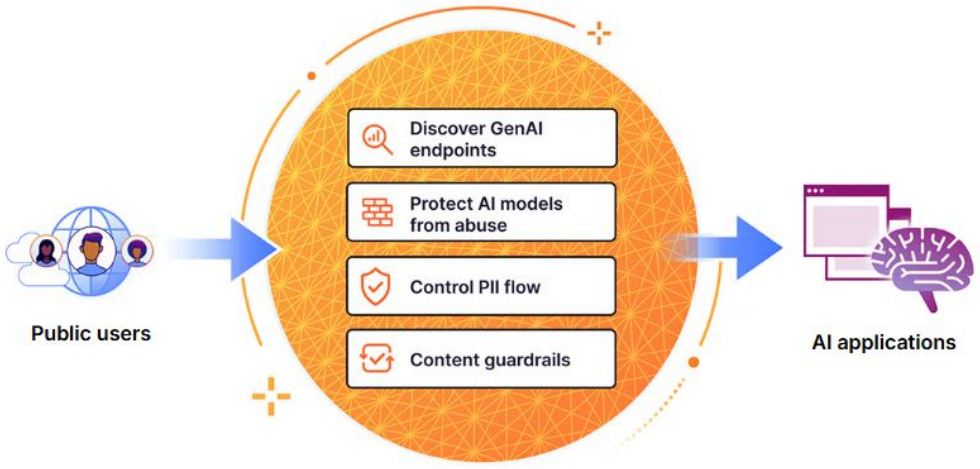

Defend AI-powered apps from cyber threats

Before public users can interact with AI applications, organizations should build an inventory of all GenAI components and deploy security controls that protect models from abuse. This ensures that that only appropriate data flows both in and out, and that models aren't "tricked" by AI exploits.

Use Case 3: Protect interactions between AI agents and internal data

An AI agent can perform tasks on behalf of human users, even without specific instructions, but they need access to much more information than just their base training data.

The Model Context Protocol emerged as a standard way for AI agents to retrieve that additional information. It connects AI agents (MCP clients) to internal resources (MCP servers) so that responses returned to users have the appropriate context. However, like any technology, MCP can be exploited by cyber attackers to leak or poison internal data sources.

There are three key recommendations to protect the critical link between AI agents and internal MCP servers:

- **Analyze activity logs**. Aggregate and review request logs to understand the interactions among AI agents and MCP servers.
- **Authenticate and authorize**. MCP supports optional OAuth-based authentication and authorization, so use it to verify, validate, and restrict access between AI agents and internal MCP servers.
- **Create an MCP server portal**. Rather than configuring MCP servers individually, connect all accessible servers with a single MCP server portal URL to simplify and unify management.

What is the Model Context Protocol?

MCP is a standard way to make new information sources and tools available to large language models (LLMs) to improve their responses.

Initially developed by Anthropic, today MCP is open source and has rapidly become an industry standard for AI agents.

Learn more

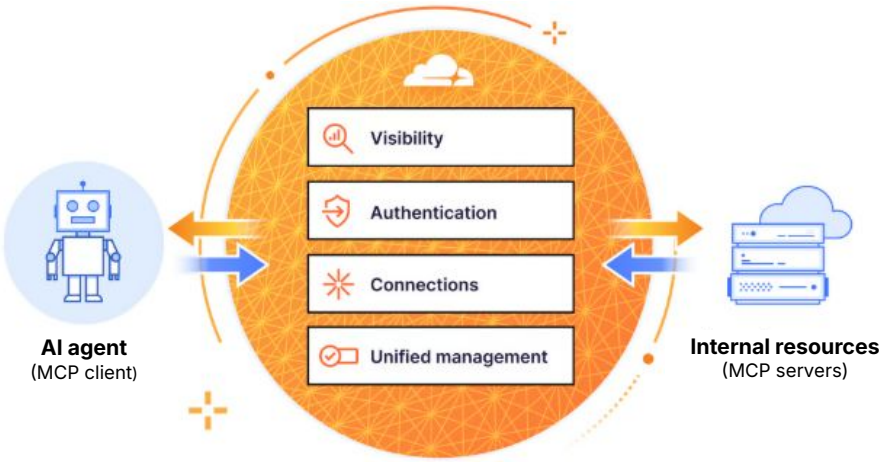

Protect interactions between AI agents and internal data

AI agents act on behalf of individuals, so it's vital to treat them them just like human identities. Appy Zero Trust principles like authentication, authorization, and least privilege to ensure trustworthy connection between AI agents an internal resources.

Use Case 4: Build innovative AI apps and digital assistants

Across the pubic sector, organizations are racing to build new AI assistants to save people time and effort. Here are a few examples:

The U.S. Citizenship and Immigration Services

"Emma" is an virtual assistant fluent in both English and Spanish to help people with immigration services, green cards, passports, and more.

The City of Atlanta

"Ava" is an AI chatbot built into the city's ATL311 app, helping people find and request non-emergency services like pothole reporting.

The State of South Carolina

"Bradley" is an AI assistant being readied to help the public to get fast assistance with taxes, water bills, and more.

However, building innovative AI apps and digital assistants are different from simplistic chatbots of the past.

Traditional (rule-based) chatbots follow fixed scripts and are sufficient for simple tasks, but they fail with "off-script" questions. AI (conversational) assistants use natural language processing, machine learning, and large language models to understand intent, handle context, learn, improve, and integrate data for comprehensive answers to complex queries. And as discussed, AI assistants must be secure from cyber attacks and malicious behavior.

Following four best practices can help public sector organizations maximize the value of these AI-powered services while minimizing the potentially serious risks.

- **Find the right use case.** Define the problem the AI assistant will solve, design for users, and iterate based on feedback and lessons learned. Focus the first iteration on a specific, common problem – like answering frequently asked questions – and learn from the experience.
- **Assemble foundational technologies.** Every AI assistant starts with a model, a data set, and a means to retrieve data. In addition to MCP mentioned earlier, Retrieval-Augmented Generation (RAG) is another method to provide models with internal context. RAG helps AI *know* more; MCP helps AI *do* more.

> **"AI chatbots can provide fast, consistent, and compliant responses — without requiring humans. But they can be difficult to build, secure, and govern."**
>
> **Dan Kent**
> Field CTO for Public Sector, Cloudflare

- **Don't rush design.** Carefully designing the interface, platform support, and relationship to other content will be key to ensuring its successful adoption. Start with user research, then define a limited scope for the chatbot, like specific topics it will cover or a certain type of question-and-answer model. Gather data and feedback to continuously improve the bot's performance and knowledge base.
- **Prioritize security and governance.** The prior use cases lead into this one, emphasizing how critical it is to build security in from the start. Understand and control the input, output, and everything in between is essential for building modern AI assistants that are truly trustworthy – an essential requirement for every public sector AI innovation.

With these best practices defined, the next step is understanding the key building blocks for AI application development:

- **Unify the AI control plane**. Manage all AI models from a single location that routes requests, caches responses, and monitors performance. It will also help save time and reduce overall costs.
- **Protect credentials at the edge**. Store API keys and secrets at the edge to prevent client-side exposure and simplify key rotation.
- **Enforce content safety guardrails**. As with Use Case 2, it's important during the build phase to automatically identify, block or redact harmful content and sensitive data in prompts and responses.
- **Secure the AI backend**. And as in Use Case 3, it's vital to connect AI agents securely to internal resources and data stores. Be sure to integrate with identity providers for authentication and authorization.

Realizing the benefits of AI requires an understanding of its risks, and these four major AI use cases provide the foundation needed to ensure that public sector AI innovations will be trustworthy and secure.

Summary

Major AI use cases

1. Control internal use of Generative AI
2. Defend AI-powered apps from cyber threats
3. Protect interactions between AI agents and internal data
4. Build innovative AI apps and digital assistants

Next steps

→ Discover more about boosting government efficiency with AI agents

→ Read the AI chatbot playbook for government

→ Learn how to confidently scale AI security

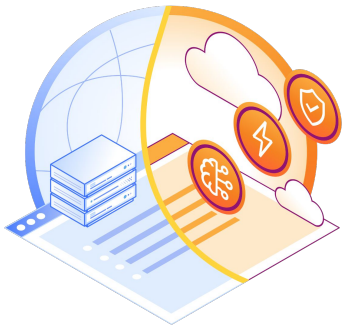

Ready to advance your AI initiative?

Visit Cloudflare for Public Sector

1 888 99 FLARE | cloudflare.com/public-sector