# The ATO Bottleneck

FedRAMP takes 18 months. AI models update every 90 days. Most federal agencies are running AI tools that have never been through a full authorization.

01

# Background

How 70+ federal agencies deployed AI ahead of the authorization process — and why the governance gap is now structural.

01 — BACKGROUND

# The ATO process was built for static systems. AI is not a static system.

- ATO (Authority to Operate): NIST SP 800-37 requires 12–18 months of security assessment before a federal agency can legally operate a new IT system. Designed for infrastructure that, once deployed, stays stable.
- FedRAMP: The federal cloud authorization program. 300+ services authorized as of 2025. Microsoft Azure Government, AWS GovCloud: FedRAMP High. Microsoft Copilot for M365 Government: FedRAMP High — one of the few AI tools with federal authorization.
- OpenAI ChatGPT Enterprise, Anthropic Claude, Google Gemini (commercial): No FedRAMP authorization. A federal employee using a personal account for government work = data spillage event.
- EO 14110 (Oct 2023): Required agencies to designate a Chief AI Officer, submit AI use case inventories, and implement NIST AI RMF practices. Most agencies are still completing the baseline.
- 70+ federal agencies with active AI use cases. DOD (CDAO), VA, GSA, IRS, CBP: all deploying AI. Authorization status: inconsistent.
- The structural gap: commercial AI models update every 60–90 days. An ATO for a specific model version is stale before the authorization process completes.

02

# Decision Required

How do you deploy AI at mission speed when the authorization process runs at procurement speed?

02 — DECISION REQUIRED

# Restrict AI to FedRAMP-authorized tools, operate under provisional authorization, or build a continuous ATO program?

FedRAMP-authorized AI tools (Copilot for Government, Azure OpenAI Service Government, AWS Bedrock GovCloud) are deployable now — but cover a subset of available models and lag commercial releases by 6–12 months.

Provisional authorization (documented risk acceptance while formal ATO proceeds) is the most common actual practice — but creates a documented governance gap that IG reviews and incident investigations can surface.

Continuous ATO (cATO) — ongoing monitoring rather than point-in-time authorization — is operationally sustainable for AI, but requires security tooling and staffing most civilian agencies have not built.

Shadow AI is the immediate risk: employees are using personal ChatGPT accounts for government work today. The authorization process does not prevent the data spillage already occurring.

03 — OPTIONS

# Four authorization postures.

Option A

**FedRAMP-authorized tools only — Copilot for Government, Azure OpenAI Gov, AWS Bedrock Gov**

Most defensible posture. Inherits existing authorizations, no independent ATO required. Limited model choice — model access lags commercial by 6–12 months. Right for agencies with low risk tolerance or rights-impacting use cases.

Option B

Recommended

**Continuous ATO (cATO) program with tiered use-case risk classification**

Operationally sustainable: ongoing monitoring replaces point-in-time authorization. Requires investment in security tooling, automation, and staffing. DOD components have moved here under DISA DevSecOps. Right for agencies deploying AI at scale.

Option C

**Provisional authorization — documented risk acceptance, formal ATO in parallel**

Most common actual practice. Legally defensible when AO genuinely owns the residual risk and monitoring requirements are defined. Risk: documented governance gap is visible in IG reviews and incident investigations.

Option D

**Internally hosted open-weight models on agency infrastructure — no commercial API dependencies**

Full change management control and authorization ownership. Operationally realistic only for DOD/IC components with dedicated AI engineering organizations. Not achievable for most civilian agencies at required scale.

04 — RECOMMENDATION

# Tier your use cases. Address shadow AI first. Brief your AO on model version changes.

Build a use-case tiering framework: Tier 1 (productivity, non-sensitive) ’ FedRAMP tools or documented risk acceptance. Tier 2 (mission support, not rights-impacting) ’ ATO scoped to AI risk surface on authorized IaaS. Tier 3 (rights/safety-impacting) ’ full ATO + OMB M-24-10 minimum practice assessment.

Address shadow AI before the formal authorization process. DLP controls on government devices + explicit policy on approved tools by data sensitivity level. The spillage risk is occurring now.

Complete the OMB AI use case inventory as a governance foundation — not a checkbox. Every use case in the inventory needs an authorization status and a rights-impact determination.

Brief your Authorizing Official explicitly on model version change management. AOs who signed AI ATOs in 2023–2024 likely do not know that model updates may have changed the system they authorized.

For Tier 2–3 use cases: establish a provider release monitoring process. Designate a model owner. Define what triggers a change impact assessment.

05 — RISKS

# Five risks every federal CIO is managing around.

1.

### Shadow AI data spillage — occurring now

Employees using personal ChatGPT accounts for government work is a data spillage event under most agency security policies. It does not wait for the authorization process. DLP + sanctioned alternatives are the only controls.

2.

### Inherited FedRAMP authorization does not cover AI failure modes

FedRAMP certifies cloud infrastructure security. It does not certify that an AI tool will not hallucinate in a procurement memo or produce a biased output in a benefits determination. The AI-specific risk gap exists across all FedRAMP-authorized AI offerings.

3.

### Model version changes create undocumented authorization gaps

Commercial AI providers update models without federal change management notifications. Every silent model update on an authorized system is a potential SP 800-37 change management deviation — accumulating undetected.

4.

### Rights-impacting AI without OMB M-24-10 assessments — enforcement increasing

Multiple agencies submitted use case inventories with rights-impacting AI (benefits, law enforcement, employment) without completed minimum practice documentation. IG reviews are now checking against this standard.

5.

### Procurement cycle misalignment — authorized tools may be obsolete before agencies use them

Federal procurement: 12–24 months. Commercial AI major releases: 6 months. By the time an agency completes procurement and authorization for a specific AI capability, the market has moved. Individual tool-by-tool ATOs cannot keep pace.

06 — QUESTIONS YOUR TEAM SHOULD BE ANSWERING

## If your team cannot answer these, that is your first deliverable.

1. Has your agency completed its AI use case inventory under EO 14110 and OMB M-24-10? For each use case: what is the authorization status, what data classification does it process, and has it been assessed against the rights-impacting AI minimum practices?
2. Which AI tools are currently in operational use at your agency without a completed ATO or inherited FedRAMP authorization? Who is the designated Authorizing Official for each, and what is the documented risk acceptance status?
3. Does your agency have a process for detecting model version changes in commercially operated AI tools — and triggering a change management assessment when the underlying model updates?
4. What is your agency's technical and policy response to shadow AI? Are DLP controls deployed? Has guidance been issued to employees specifying which AI tools are approved for which data sensitivity levels?
5. For AI tools in rights- or safety-impacting use cases: has a formal OMB M-24-10 minimum practice assessment been completed? If not, which tools are unassessed and what is the remediation timeline?
6. Has your Authorizing Official been briefed specifically on the model version change management gap — that commercial AI updates may not trigger the change management assessment SP 800-37 requires?

AI INSIGHT LAB

# The Deployment Memo

One enterprise AI deployment, dissected every Tuesday.
Written for executives who have to decide, not just read.

Subscribe at aiinsightlab.cloud — free during beta.