

# Data isn't enough: Government AI requires enterprise knowledge





## Introduction: Government AI is operating without a complete picture

Government agencies are investing aggressively in AI—layering generative AI, copilots, and advanced analytics on top of modern data platforms and mission systems.

On the surface, the architecture looks complete.

But underneath, it tells a different story.

Today's environment resembles a complex control room with dozens of disconnected screens—each representing a system: case management, email, shared drives, content repositories, data platforms, operational applications, and mission-specific systems. Every screen contains part of the mission picture, but none show it in full.

### Critical information is scattered:

- Documents and case files sit in content repositories
- Data lives in platforms and warehouses
- Transactions live in business systems
- Context is buried in emails, reports, narratives, policies, and records

There is no single, continuous view of truth—only fragments.

AI is now being deployed into this environment. But instead of seeing the full picture, AI is effectively looking through narrow windows into isolated systems, pulling partial, inconsistent, and often unverified information.

### The result is predictable:

- AI outputs lack context and reliability
- Hallucinations increase because key information is missing
- Users don't trust the results
- Mission outcomes stall at the pilot stage

In essence, agencies are trying to run AI on a foundation where knowledge is fragmented, not activated. What's missing isn't more data—or even better models. It's the ability to connect, govern, and deliver knowledge across systems as a complete, operational picture.

Until that gap is closed, AI will continue to operate in the dark—no matter how advanced the technology becomes.



## The situation: AI investment is accelerating—but outcomes are lagging

Government agencies are making significant investments across data platforms, operational systems, and AI capabilities. These investments are improving access to data, but they are not translating into consistent, mission-level outcomes. Why? AI is being deployed without complete, trusted context.

Whether your mission involves tax collection, grants administration, or law enforcement, executing its mission requires a combination of human-generated content, machine-generated data, and transactional data.

Together, these form the full context required for mission success.

The challenge is that in most agencies, these types of information are typically managed separately across independently managed systems, with varying levels of structure and governance.

Without bringing them together in a governed way, both AI and mission execution operate in an incomplete and often unverified context. When AI is deployed against fragmented information, it can increase risk, reduce trust, and limit impact.

## The critical shift: From data management to knowledge activation

There's been a fundamental shift in how agencies need to think about unstructured content.

Traditionally, unstructured content, mainly comprised of images, audio, and video files, was managed for storage, archival, records, and compliance requirements. Basic retrieval when needed was the goal. It was fundamentally a back-office function.

But here's what's changed: unstructured content is now a primary input to AI and mission execution. With the rise of generative AI and advanced analytics, content that was historically managed for compliance is now foundational to how agencies operate.

This means unstructured content must be analyzed, enriched with metadata and context, governed with security and access controls, and made securely accessible across systems. It's no longer enough to just store it—agencies need to activate it and make it usable in context.

The key implication is this: AI effectiveness and mission outcomes now depend on the ability to operate on enterprise knowledge—by activating unstructured content as AI-ready, governed context.

This isn't a content management problem anymore; it's a strategic capability for AI readiness and mission execution. And it's a capability most agencies don't yet have in place.

## The opportunity: A governed knowledge activation layer

To unlock AI at mission scale, agencies need a governed knowledge activation layer.

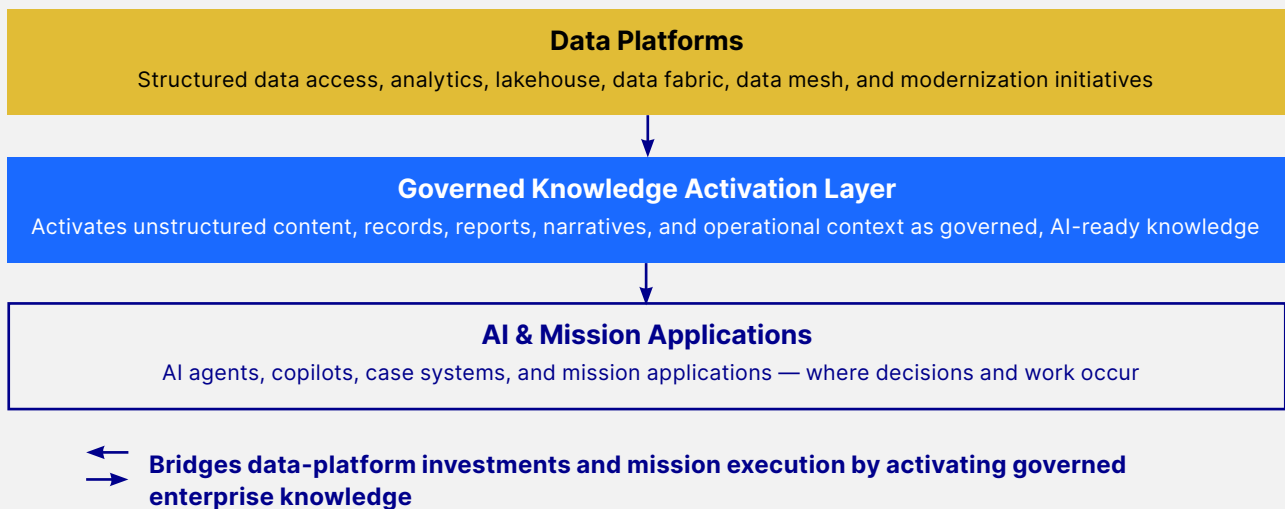
This is not a product. It is an architectural capability layer within the agency's information environment—one that connects, enriches, governs, and delivers trusted enterprise knowledge across data platforms, content repositories, mission systems, operational applications, and AI tools.

A governed knowledge activation layer helps agencies:

- Connect information across existing repositories, platforms, and applications.
- Enrich content and data with metadata, entities, relationships, and context.
- Preserve governance, permissions, provenance, records, and auditability.
- Deliver trusted knowledge into AI agents, search, analytics, copilots, workflows, and mission systems.

### A governed Knowledge Activation Layer

An architectural capability layer that activates trusted enterprise knowledge for AI and mission outcomes



The goal is not wholesale content migration or forcing all information into a single repository. In many cases, information can remain where it already resides. But some high-value content, such as official records, case artifacts, evidentiary documents, decisions, or mission-critical content, may need to be selectively enriched, copied, archived, promoted, or managed within a governed content platform when stronger lifecycle control, auditability, compliance, or workflow integration is required.



OpenText helps agencies implement and operationalize this capability through OpenText Knowledge Discovery and related information management, governance, integration, and AI capabilities.

OpenText can federate discovery across content repositories, enrich information with metadata and relationships, apply governance and security at the object level, and connect structured data with unstructured content into usable, contextual knowledge.

A governed knowledge activation layer does not replace existing data platforms, mission systems, content repositories, or AI tools. It helps activate the knowledge they depend on and enables existing investments to work together as a system.

As agencies mature, this same governed knowledge foundation can also support agentic workflows, where tools such as OpenText Aviator Studio help create and orchestrate agents that act on trusted enterprise knowledge within mission processes.

## **The benefits of a governed knowledge activation layer**

Once you have a knowledge activation layer in place:

### **It reduces AI hallucinations and builds trust.**

When AI agents are grounded in governed, verifiable context from within, not just generic training data, errors decrease significantly. More importantly, government teams begin to trust the outputs because they can trace them back to authoritative sources with clear provenance.

### **It accelerates agent deployment across platforms.**

Because knowledge is activated and governed through a shared architectural layer, it can be delivered into multiple environments, including Salesforce, Microsoft, ServiceNow, SAP, Oracle, custom mission applications, and open-source frameworks.

Agencies do not need to rebuild context or recreate governance for each platform. They can enable consistent, governed knowledge wherever work happens.

### **It unlocks institutional knowledge at scale.**

Government agencies have accumulated decades of enterprise content, case files, transactions, correspondence, records, and operational data. A governed knowledge activation layer helps transform that information into searchable, usable, contextual knowledge for AI and mission users.

In some cases, that knowledge can remain federated from its existing source. In others, selected content may be brought under stronger governance to support records management, case execution, evidentiary control, auditability, or workflow needs.

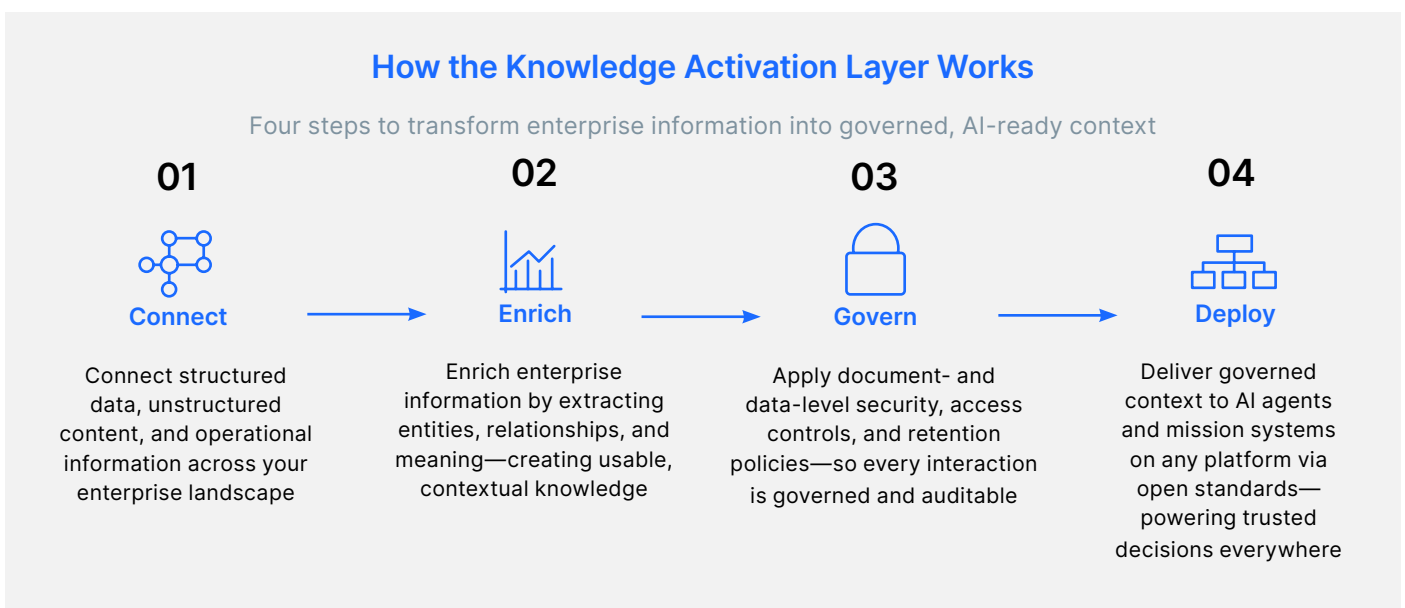
## It ensures compliance from day one.

Governance, access controls, provenance, auditability, and retention must be designed in—not added later.

Every retrieval should be governed. Every interaction should be traceable. Policies should be enforced consistently across users, systems, and AI agents.

This approach applies across government use cases, including case management, investigations, grants management, benefits delivery, legal and regulatory compliance, emergency response, procurement, HR, and transportation safety.

## How it works



**Step one: Connect.** OpenText connects structured data, unstructured content, and operational information from across an agency's technology landscape. This includes data lakes, content repositories, transactional systems, and business networks.

**Step two: Enrich.** OpenText enriches enterprise information by extracting entities, identifying relationships, and building usable context. This is where data products and knowledge graphs are created—capturing entities, relationships, and meaning. This is how information becomes governed, usable context for AI and mission execution.

**Step three: Govern.** OpenText applies document-level security, access controls, and retention policies. Every retrieval is governed and auditable. If someone shouldn't have access to a document in SharePoint, they won't have access to it through AI either. Governance follows the data.

**Step four: Deploy.** Finally, OpenText exposes this governed context to AI agents on any platform via open standards. Whether an agency is using MCP, vector databases, RAG architectures, or agent frameworks, we integrate with any approach.



This is how raw information becomes governed, usable context for AI and mission execution.

In each of these cases, the difference isn't just AI—it's that the AI is operating on **complete, governed context drawn from across systems**.

## How OpenText supports the knowledge activation layer

OpenText brings decades of experience helping public sector and regulated organizations manage, govern, secure, enrich, and operationalize high-value enterprise information.

That experience matters because trusted AI depends on more than access to data. It requires governed knowledge drawn from content, records, operational systems, transactional information, and mission workflows, with security, provenance, and compliance preserved throughout.

OpenText helps agencies implement and operationalize a governed knowledge activation layer through capabilities that:

### Connect across existing systems

OpenText works across content repositories, data platforms, Microsoft, Salesforce, SAP, ServiceNow, and mission-specific applications, helping agencies activate knowledge without forcing wholesale system replacement.

### Govern information at enterprise scale

OpenText applies security, access controls, retention, auditability, records management, and compliance to information throughout its lifecycle.

### Activate content, data, and context for AI

OpenText helps transform distributed information into enriched, AI-ready knowledge by extracting entities, identifying relationships, applying metadata, and delivering governed context into AI, search, analytics, and operational workflows.

### Support mission-critical public sector environments

With more than 3,500 public sector customers globally, OpenText understands the operational, security, compliance, and modernization constraints agencies face.

Together, these capabilities help agencies move from fragmented information to governed enterprise knowledge—the foundation trusted AI requires. The key here is that this works with any AI agent platform. Agencies are not locked in. A knowledge activation layer should serve the public sector's entire AI strategy, not just one vendor's tools.



## Bottom line

AI in government will not succeed on data alone. It requires complete, governed, operational knowledge. Without it, AI remains experimental. With it, AI becomes mission critical.

## What this enables

- Reduced AI hallucinations
- Accelerated AI deployment
- Unlocked institutional knowledge
- Built-in compliance and governance
- Better mission decisions across systems

## Next steps

- 1. Review your mission context.** We are prepared to review your agency's specific data platform initiatives, your AI goals, and your operational challenges. Every agency has unique requirements. Let's talk about yours.
- 2. Explore reference architecture.** Let us show you specifically how OpenText integrates with your existing data platforms and systems. We'll map this to your architecture, not a generic diagram.
- 3. Identify entry opportunities.** Let's define specific tactical wins: AI accuracy improvements, search modernization, knowledge access projects—that build toward your strategic goals. How can OpenText create mission value over the next 90 days?

[Download Enterprise Artificial Intelligence: Building Trusted AI in the Sovereign Cloud >](#)

[Schedule a demo >](#)



## Why OpenText?

- 1. OpenText is trusted.** We have more than 3,500 enterprise public sector customers globally. We have decades of experience as a trusted data custodian in regulated industries, from banking to healthcare to government. We understand compliance, security, and mission-critical operations.
- 2. OpenText is an expert in transactional intelligence.** We operate the largest business network of its kind with 31 billion B2B transactions annually, connecting more than one million trading partners. That transactional data provides intelligence that few vendors can match.
- 3. OpenText is unique in combining tri-domain information**—that is, human-generated content, machine-generated operational data, and transactional data—into a single governed layer. Most vendors handle one or maybe two of these. We handle all three because mission work requires all three.
- 4. OpenText works with everyone.** Whether your agency is building on Microsoft, Salesforce, SAP, ServiceNow, or open-source frameworks, your knowledge layer must work with all of them. This is strategic: your knowledge infrastructure should outlast any single vendor's AI platform.

All this adds up to public sector trust in OpenText to support mission-critical operations and secure data management. We understand your requirements, your compliance needs, and your operational constraints.

## Contact



### Keith Nelson

Senior Strategist for Public Sector,  
OpenText

[nelsonk@opentext.com](mailto:nelsonk@opentext.com)

The Hon. Keith Nelson is Industry Senior Strategist for Public Sector at OpenText. Mr. Nelson has spent more than 20 years working in the fields of public sector consulting and high-tech marketing and as a government appointee. His roles in government include serving as Assistant Secretary for Administration (a position requiring Senate confirmation), Acting Assistant Secretary for Research and Technology, Acting Chief Financial Officer, and Deputy Chief Information Officer at the U.S. Departments of Transportation, Labor and Housing and Urban Development.