

AI governance is becoming an operating model

Insights from the September 2026 Kyndryl US Government CIO Expert Exchange

kyndryl®



Contents

01

Overview

02

Measure value

03

Buy versus build

04

**Governance that
enables**

05

Data gates scale

06

**US Government CIO
Expert Exchange**

01

Overview

The exchange discussed how AI should be governed, funded, measured, and sustained. Leaders favor room for early experimentation, but draw a firmer line around production use, recurring cost, enterprise data, and operational accountability. The strongest approaches combine accessible tools, practical training, communities of practice, smaller groups of power users, documented standards, and a visible route for moving promising use cases into secure environments.

The conversation also makes clear that AI strategy cannot be separated from core technology management. Platform economics, application ownership, data classification, privacy obligations, vendor leverage, workforce adoption, and digital identity all shape whether a promising prototype can become a dependable public service.

Explore the 2026 ISG Provider Lens:
Public Sector Services and Solutions report



02

Measure what government actually values

Traditional ROI remains important, especially when an AI-enabled solution avoids a planned expenditure or replaces a costly contract. The exchange cites examples where in-house development avoided significant external spend, including a \$500,000 product quote and another initiative that generated \$1.6 million in savings. Leaders view these direct dollars as the most defensible financial evidence for budget and oversight conversations.

Saved hours do not automatically translate into budget reductions when staffing levels remain unchanged; instead, they create capacity. Their value often appears as faster processing, improved service delivery, stronger compliance, greater transparency, or the ability to redirect staff to higher-value work. One ballot-review use case reportedly removed roughly four review rounds and about 2,000 hours of review time while retaining human review at the decision point.

AI Maturity Assessment

Use a dual value framework: pair cost-benefit measures with public value outcomes, and establish baseline time and error rates before automation begins.

- Define the intended outcome before committing resources to a modernization use case.
- Distinguish avoided spend from capacity released or service improved.
- Measure quality and error reduction alongside speed.



03

Buy versus build is now a strategic lever

AI-assisted development changes the economics of the traditional buy-versus-build decision. Small teams can prototype targeted capabilities quickly, exposing the gap between enterprise software pricing and the narrower functionality an individual agency actually needs. This gives public-sector buyers new leverage in vendor negotiations and creates credible alternatives when renewal or replacement costs become difficult to justify.

The trade off is operating responsibility. A prototype built by one developer may be compelling, but it can also create ownership, security, maintenance, documentation, and succession risks. The practical dividing line emerging from the exchange is to exhaust embedded capabilities in major platforms for native capabilities, while considering custom development when a specialized need creates material value that the platform cannot deliver economically.

- Use rapid prototypes to test vendor claims and strengthen negotiations.
- Set an explicit threshold for when custom development is economically justified.
- Require a transition plan from individual prototype to agency-owned service.

“The real decision is not just buy or build, but whether you can securely own, operate, and retire what is built.”

Government CIO Expert Exchange participant



04

Guardrails must create a path to yes

Shadow AI grows when employees can see the opportunity but cannot access an approved way to act. One agency scan identified approximately 1,200 users among roughly 4,500 to 4,800 office employees using public AI tools, while interns and business units were also creating applications before central IT had visibility. Leaders warn that blanket restrictions can push experimentation onto personal devices or into environments with even less oversight.

More effective governance pairs approved tools with clear standards, security training, use-case intake, and accessible review processes. Communities of practice and internal showcases help identify champions, spread practical knowledge, and bring emerging solutions into view. As citizen development expands, central technology teams increasingly become commissioners and stewards: they establish the boundaries, assess scalability, and determine when a locally created tool should become an agency-wide capability.

Policy-As-Code

"Governance should make the secure route easier to find and use than the workaround."

Government CIO Expert Exchange participant

- Provide approved tools and simple guidance before imposing broad restrictions.
- Use communities of practice to surface innovation and reinforce standards.
- Track agent ownership and create a process for transferring successful tools to enterprise stewardship.



05

Data access is the constraint on scale

The exchange consistently separates the AI engine from the data required to make it useful. Teams can build capable agents quickly, but enterprise use stalls when those agents need protected student, resident, financial, operational, or cross-departmental information. Privacy obligations and distributed data ownership therefore become more consequential than model availability.

This challenge requires a joint operating approach across technology, legal, privacy, security, and business leadership. One proposed monitoring practice uses unusual data-movement volume as a trigger for further review without inspecting the underlying content. The broader objective is to create governed access patterns that let approved applications use the minimum necessary data while preserving visibility and accountability.

AI maturity will increasingly be defined by the organization's ability to provide trusted, classified, and permissioned data, not by access to another model.



- Map data owners and legal constraints before scaling high-value use cases.
- Create reusable access patterns and classification rules for common AI scenarios.
- Monitor agent activity and data movement as part of enterprise security operations.

06

What is the Kyndryl US Government CIO Expert Exchange?

The Kyndryl US Government CIO Expert Exchange is an exclusive peer-to-peer community of senior IT leaders, CIOs and CTOs who gather quarterly for open virtual discussions to share their top-of-mind challenges and gain insights from their peers.

The topics for the roundtables are determined based on interests and feedback provided during 15-minute prep calls conducted with participants in the weeks leading up to the roundtable.

The September 2026 conversation was hosted by:



Jeff Pany
*Enterprise
Architect,
Kyndryl*





© Copyright Kyndryl Inc. 2026. All rights reserved.

Kyndryl is a trademark or registered trademark of Kyndryl, Inc. in the United States and/or other countries. Other product and service names may be trademarks of Kyndryl, Inc. or other companies.

This document is current as of the initial date of publication and may be changed by Kyndryl at any time without notice. Not all offerings are available in every country in which Kyndryl operates. Kyndryl products and services are warranted according to the terms and conditions of the agreements under which they are provided.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions. Kyndryl products and services are warranted according to the terms and conditions of the agreements under which they are provided.