KUMIHO STRATEGIES
OPERATIONAL PERSPECTIVES SERIES

Operational Throughput Without Operational Drift

Trust, Governance, and Visibility in AI-Accelerated Federal Modernization

Kumiho Strategies | 2026

A Note on This Document

This paper was researched, synthesized, and produced using AI-assisted workflows. That fact is not incidental — it is part of the argument. The analysis, judgment, framing, and conclusions are human. The acceleration is not. That distinction is precisely what this paper is about.

EXECUTIVE SUMMARY

The Transition Is Already Underway

Federal agencies and their contractor partners are already operating inside an AI-accelerated environment. The throughput is real. The speed is real. What is not keeping pace is the governance, accountability, and trust infrastructure that makes accelerated systems coherent, legitimate, and sustainable.

This is not an argument against modernization. Many federal missions cannot afford to slow down. What this paper argues is more specific: acceleration without governance maturity creates operational drift — and drift has a cost. For delivery. For workforce trust. For public legitimacy. For institutional viability.

> ***The organizations that will define responsible AI-assisted delivery are not the ones waiting to see how this plays out. They are the ones building the governance infrastructure now, while everyone else is still figuring out the tools.***

This pattern is not new. Cloud First taught the federal enterprise what happens when technology adoption outpaces governance discipline. AI is amplifying that lesson at higher speed, with less visibility, and with less institutional memory of what went wrong the last time.

The pages that follow examine how this transition is already unfolding — where the real gains are, where the governance debt is compounding, what happens to public trust when systems become invisible, and what responsible operational posture looks like for the organizations that want to lead rather than react.

SECTION I

We Have Seen This Before

The Cloud First Warning and the Data Governance Gap

Before treating AI acceleration as an entirely new problem, it is worth pausing on a pattern the federal government has navigated before — because the lessons from that experience apply directly to the moment we are in now.

Cloud First was a legitimate modernization priority. The operational arguments were sound: scalability, resilience, cost efficiency, and the retirement of aging infrastructure that had become expensive and fragile. Agencies moved. Programs migrated. Budgets shifted.

But a critical distinction was consistently lost in the urgency of transition: migration is not modernization.

Moving data to a cloud environment does not automatically make that data well-governed, interoperable, or useful. Many agencies lifted fragmented systems and placed them into scalable environments — and then discovered that scale amplifies fragmentation rather than resolving it.

> ***The hardest part of modernization was never the technology. It was the governance, the accountability, and the coherence required to make new capability sustainable. That was true for cloud. It is more true for AI.***

The result, across much of the federal enterprise, is an environment that practitioners recognize: data that technically exists in accessible systems but cannot be reliably used. Inconsistent metadata. Unclear ownership. Duplicated repositories. Incompatible taxonomies. Governance structures that were never designed for the interoperability expectations now being placed on them.

This was not a failure of cloud technology. It was a failure of governance maturity to keep pace with adoption speed.

AI is now encountering this same environment — and the consequences are compounding. AI systems are heavily dependent on data quality, structure, consistency, and governance discipline.

When they operate on fragmented, poorly governed data ecosystems, they do not compensate for that fragmentation. They amplify it. Outputs inherit the inconsistencies of the inputs. Synthesis reflects the incoherence of the underlying sources.

Ask yourself whether your organization can answer these questions today: Who owns your data, and who is accountable for its quality? If an AI-assisted output is wrong, how quickly can you trace it back? If an auditor asked to walk through a decision your systems made last quarter, could you show them the thread?

If the answer is uncertain, AI adoption is not solving your governance problem. It is accelerating it.

This is why AI readiness is increasingly a governance maturity question, not primarily a technology question. Organizations that invested in data discipline, ownership clarity, and workflow coherence during the cloud era are better positioned to extract genuine value from AI. Organizations that treated migration as modernization are discovering that AI reveals the debt they accumulated — faster.

The pattern is consistent enough to be called a behavior, not a mistake.

SECTION II

The Acceleration Era

What Is Already Happening in Federal Delivery

The federal operating environment has shifted substantially in the span of a few years. Not because a single policy created the change, but because several pressures arrived simultaneously and began reinforcing each other.

The gains are real and worth naming.

NOAA deployed a new suite of AI-driven global weather prediction models in December 2025. The AIGFS system delivers a full 16-day forecast using just 0.3% of the computing resources previously required, finishing in approximately 40 minutes where traditional models took hours.[1]

NASA's Perseverance rover has completed more than 90% of its Mars surface travel autonomously — a proof point for AI-assisted decision-making in high-stakes, low-oversight environments that has direct implications for how autonomous systems will be governed across the federal enterprise.[2]

The Department of State launched StateChat — its enterprise AI chatbot approved for Sensitive but Unclassified data — in August 2024. In its first week alone, more than 1,900 employees used it for the first time and it received more than 10,000 prompts in a single day. Department-wide use has grown exponentially since, supporting translation, drafting, summarization, and analysis across State personnel globally.[3]

Across agencies surveyed by GAO, documented AI use cases nearly doubled between 2023 and 2024. The 2025 federal inventory now exceeds 3,600 reported use cases — a 69% increase over 2024, and roughly five times the number reported in 2023.[4]

These are not pilot programs or press releases. They are operational realities that are changing what agencies and their delivery partners can reasonably be expected to produce — and how fast.

This document is itself a case study in that shift. The research synthesis, structural architecture, and drafting that produced this paper were developed through AI-assisted production in a fraction

of the time a traditional cycle would require. The analysis, the framing, the operational judgments, and the conclusions are human. The speed is not. That distinction matters — and the transparency about it matters more.

> ***AI is not coming to federal operations. It is already there. The question is not whether to engage it. The question is whether your organization is governing it or just using it.***

OneGov — launched by GSA in April 2025 and now at its one-year mark — represents a different kind of structural pressure. In its first year, GSA executed 20 unified agreements with major technology suppliers and has reported $1.1 billion in savings through negotiated discounts. At the same time, reporting at the one-year mark describes agency IT leaders reflecting with "praise, but also confusion over how the program works" — a candid signal that the initiative is delivering results faster than organizations can fully absorb and operationalize them.[5]

What OneGov is already changing is the expectation. Agencies can no longer treat their data, workflows, and governance structures as purely internal concerns. Cross-agency visibility, standardized exchange, and interoperable delivery are becoming the baseline assumptions — not future aspirations. For contractors and delivery partners, that shift has direct implications for how they structure accountability, govern tooling, and represent their work to oversight bodies.

The organizations that understand this now — before it is formalized in solicitations and source selection criteria — will not have to scramble to explain themselves later.

SECTION III

Visibility, Trust, and Legitimacy

What Happens When Systems Become Invisible

In 2025, the New Zealand government made a decision that would have been unthinkable a decade ago: it cancelled the 2028 traditional census and announced it would replace it with administrative data drawn from existing government records — tax filings, health databases, driver licenses, and social service records — combined with annual sample surveys.

The operational logic was straightforward. The 2023 census cost NZ$320 million — more than double the per-capita cost of a decade prior — while still falling short of its 90% response rate target. Stats NZ noted publicly that it was becoming harder to motivate people to complete census forms, and that some were increasingly hesitant to share personal information with government.

The public and expert response was instructive — and immediate.

Civil liberties organizations submitted formal objections, noting that Stats NZ had structured its consultation to make the substantive change appear inevitable and feedback "off topic." A government-appointed independent expert panel recommended against moving to the new model by 2028, citing significant gaps in existing administrative data and concerns about quality. Māori communities raised particular concerns about under-representation and Indigenous data sovereignty. Former Government Statistician Len Cook called publicly for independent review. None was granted.[6]

> ***People do not distrust AI in isolation. They distrust opacity, invisible decision-making, and the loss of agency over systems that govern them. That is not an AI problem. It is a legitimacy problem.***

The traditional census carried something beyond data collection. It was a visible, participatory act — a recognizable ritual in which citizens and government exchanged information through an explainable process. People knew what they were providing. The form on the doorstep was not just logistically necessary. It was institutionally legible.

Administrative-record-driven systems replace visible participation with invisible integration. The government already knows. The data is already collected. The count is being assembled from systems most people do not know they are interacting with in this context.

Even where this is entirely legal, accurate, and well-intentioned, it changes the psychological relationship between the institution and the person being governed. Consent perception changes. Transparency perception changes. Sense of agency changes. And when institutional trust is already fragile — as it was in New Zealand following the COVID era — invisible systems accelerate that erosion.

The United States is likely to face identical pressure within the decade. The Census Bureau has explored administrative data integration to reduce the cost and burden of decennial operations. The operational arguments are the same. So are the trust dynamics.

This dynamic is not limited to census operations. It is a characteristic of AI-accelerated governance environments broadly. When AI-assisted systems make eligibility decisions, when fraud analytics flag cases for review, when automated intake processes route requests without visible human touch, when benefits are adjudicated through algorithmic systems — the same legitimacy questions emerge. The workforce experiences them too.

The governance gap between AI deployment and accountability infrastructure is not theoretical — it is documented.

The 2024 OMB Federal AI Use Case Inventory recorded 2,133 AI use cases across the government. Of those, 351 were flagged as rights-impacting, safety-impacting, or both. Analysis of the primary inventory data shows 125 of those systems requested compliance extensions at time of reporting. OMB's December 2024 published release noted 206 extensions granted at publication — a figure that shifted as agencies updated submissions through 2025. Either number tells the same story: a substantial portion of rights- and safety-impacting AI systems were operating without certified safeguards in place.[7]

The agencies carrying the largest share of that burden are not the ones operating carelessly. They are the ones operating at the hardest intersection of AI capability and human consequence — under mission pressure that does not pause while governance frameworks catch up.

The Department of Veterans Affairs, for example, is running AI in environments where the stakes could not be higher: medical imaging, wound assessment, cardiac monitoring, diagnostic support for clinicians serving veterans. Of 145 systems flagged as rights or safety impacting, 109

requested compliance extensions. That is not evidence of negligence. It is evidence of an agency deploying AI at mission speed inside a compliance architecture that was not built for the pace of adoption now being asked of it. The VA knows this. Its teams are working the problem. The gap is systemic, not organizational.

The Social Security Administration presents a similar picture. Its AI systems are embedded in the most consequential public-facing processes in the federal government — disability determinations, benefits adjudication, fraud detection, overpayment review. Every one of its rights or safety impacting systems requested a compliance extension. The fact that those systems are being inventoried and reported at all reflects an agency that is engaging the accountability question honestly, even when the honest answer is that the framework hasn't caught up to the operational reality.

This is what the governance gap actually looks like in practice. Not bad actors. Not indifferent organizations. Capable agencies operating at the edge of what AI can do, inside accountability structures that the speed of adoption has outrun.

The contractor community is not exempt from this dynamic either.

Between November 13, 2025 and May 15, 2026 — approximately six months — a contractor working for CISA, the federal government's primary civilian cybersecurity agency, maintained a publicly accessible GitHub repository named "Private-CISA" containing 844 megabytes of sensitive data: plaintext passwords, AWS GovCloud administrative credentials, SAML certificates, and authentication tokens for internal CISA systems. The contractor had disabled GitHub's built-in secret scanning. GitGuardian's automated systems sent nine alerts to the account holder over two months; none received a human response. A security researcher who ultimately discovered and disclosed the exposure called it the worst leak he had witnessed in his career.[8]

The agency responsible for the federal government's cybersecurity posture had its own access keys sitting in an unlocked public lobby for six months — not through malice, but through a failure of basic governance discipline around everyday tooling, compounded by an absence of oversight capable of catching it.

That is operational drift made visible. And it is a more honest picture of where the governance gap actually lives than any policy document suggests.

SECTION IV

The Governance Maturity Gap

Five Dimensions of Operational Trust

The primary challenge in most federal modernization environments is not technological capability. It is governance maturity. And AI does not resolve governance immaturity — it exposes it faster.

Governance maturity in operational terms is not primarily about policy documents or compliance frameworks, though those matter. It is about whether an organization can answer a set of basic questions with confidence:

- Who owns this data, and who is accountable for its quality?
- When an output is produced, who validated it, and by what standard?
- If something goes wrong, where does accountability sit, and how is it traced?
- Can an oversight body follow the logic from input to output in a way that satisfies audit requirements?
- Do the people delivering work understand where AI contributed and where human judgment was applied?

These questions are not new. They existed before AI. What AI does is accelerate the rate at which they must be answered — and reduce the tolerance for ambiguity in the answers. When work happens faster, governance gaps compound faster. When outputs are produced at higher volume, accountability structures are tested at higher frequency.

Most organizations are not yet answering these questions consistently. The ones that get there first will not just be safer — they will be more competitive.

KUMIHO STRATEGIES · OPERATIONAL PERSPECTIVES

The Operational Trust Framework

Five dimensions organizations must actively maintain as AI accelerates delivery and systems become less visible.

1 Visibility

Stakeholders — internal and external — can observe what is happening in the workflow. The use of AI tools is disclosed and understood by those with oversight responsibility.

> *Can a client or oversight body see, in plain terms, where AI was used in producing this work?*

2 Accountability

There is a named human responsible for every significant output, decision, and judgment. Accountability chains are clear and auditable — not diffuse and implied.

> *If something goes wrong, who is accountable, and how quickly can that be established?*

3 Explainability

The organization can explain, in plain language, how a decision was reached, what information it was based on, and what role AI played in its production.

> *Can this decision be walked back, step by step, in a way that satisfies an auditor or a court?*

4 Ownership

Data ownership is clearly assigned. Outputs are clearly attributed. There are legible boundaries between what AI generated and what humans validated.

> *Who owns this data, and who is accountable for its quality?*

5 Human Validation

There is a consistent, documented process by which humans review and take responsibility for AI-assisted outputs before they are acted upon. Validation standards do not change because AI was involved.

> *What is the review process, and does it hold regardless of how the output was produced?*

"AI readiness is governance readiness. The organizations that understand this early will shape what responsible AI-assisted delivery looks like for the rest of the market."

KUMIHO STRATEGIES

Organizations that invest in these five dimensions are not just managing AI risk. They are building the operational infrastructure that will differentiate them as AI becomes universal and oversight expectations harden.

The organizations that treat governance as an afterthought will find that AI amplifies their operational weaknesses in ways that become increasingly difficult to manage — and increasingly visible to the clients, partners, and oversight bodies they work with.

For federal contractors, governance maturity is increasingly a client-facing value proposition. The ability to demonstrate it — not just claim it — is what oversight-conscious federal clients need from delivery partners. The questions are already beginning to appear in solicitations, in source selection criteria, and in conversations that shape teaming decisions before a solicitation ever drops.

SECTION V

Toward an Operational Transparency Model

What Responsible Disclosure Looks Like in Practice

One of the most practical things that organizations can do right now — agencies, contractors, advisory firms — is establish and disclose a clear operational posture around AI usage. Not because regulation currently requires it in most contexts. Because transparency builds trust, and trust is an operational variable.

The absence of disclosure norms is creating a gap that will eventually be closed — by oversight, by incident, or by market expectation. Organizations that establish transparent postures voluntarily will be ahead of that curve. Organizations that wait will be reacting under pressure, in ways that look defensive rather than principled.

What follows is a plain-language operational transparency framework. Not a compliance standard. Not a regulatory template. A starting point for organizations that want to communicate responsibly about how AI is integrated into their work — and that want their clients, partners, and workforces to be able to trust the answer.

AI Usage Transparency: An Operational Disclosure Model
How we use AI: AI tools assist with research synthesis, document drafting, structural analysis, and iterative refinement. AI is used as a production accelerator, not a decision-maker.
Where humans remain accountable: All strategic judgments, factual conclusions, recommendations, and final work products are reviewed, validated, and owned by human practitioners. AI outputs are treated as drafts requiring human review.
How data is handled: Client data, sensitive information, and proprietary materials are not submitted to AI tools without explicit authorization. Data handling follows established confidentiality and security protocols.
What is not automated: Analysis requiring contextual expertise, ethical judgment, strategic decision-making, and client-specific knowledge remains human-led. AI does not make operational decisions independently.

How outputs are validated: AI-assisted outputs go through the same review and quality assurance processes as traditionally produced work. Validation standards do not change because AI was involved in production.

How this will evolve: AI tooling is changing rapidly. Our practices evolve alongside the technology. We commit to maintaining transparency about significant changes in how we integrate AI into our work.

This framework is intended to be adapted, not adopted wholesale. Different organizations carry different risk profiles, different client expectations, different regulatory environments. But the underlying principle is consistent: transparency about AI usage is not a marketing exercise. It is a governance practice.

The AI governance conversation in federal environments has been heavily concentrated on PII protection and data security. Both are legitimate and necessary. But they address the surface of the problem, not its depth. The harder questions — about accountability chains, explainability, operational legitimacy, and workforce trust — are less frequently engaged and more consequential over time.

Organizations that take them seriously now are not just managing risk. They are building the kind of credibility that outlasts any single contract, any single administration, and any single tool.

CLOSING

The Next Phase of Modernization

The federal environment is navigating AI acceleration at scale — across agencies with different maturity levels, different governance cultures, different workforce readiness, and different relationships with the public trust that ultimately makes institutional operations legitimate. OneGov is adding interoperability demands on top of existing governance debt. Administrative data systems are quietly replacing visible participation rituals in ways the public and the workforce are only beginning to register.

Most organizations are still treating acceleration and governance as separate conversations. A Brookings Institution analysis of the 2025 federal AI inventory found that more than 85% of high-impact deployed AI use cases lacked some required accountability information. The technology conversation is moving faster than the governance conversation in almost every federal environment surveyed. That gap is not narrowing on its own.[9]

That separation is the problem.

> ***AI readiness is governance readiness. These are not parallel tracks. They are the same track. Organizations that are running fast on one and walking on the other are not modernizing. They are accumulating risk at speed.***

The firms and agencies that navigate this transition successfully will not be the ones with the most sophisticated AI platforms. They will be the ones that can maintain operational coherence, institutional trust, and human accountability while work accelerates and systems become less visible.

The acceleration is real. The operational gains are real. The governance infrastructure that makes those gains sustainable — in most organizations — is not yet real. That gap is the work. And it is not waiting.

Most organizations operating in this environment already know that. The governance conversation is not happening because no one understands the problem. It is not happening at the pace the environment requires because knowing and acting are separated by procurement

cycles, budget constraints, competing priorities, and the reasonable human tendency to manage the fire that is already burning before addressing the one that hasn't started yet. That is not a character failure. It is an organizational pattern. And it is one that the market — through oversight requirements, solicitation language, and client scrutiny — is beginning to price.

The institutions that close the gap first will not just be better positioned. They will have earned the right to be trusted with what comes next. That distinction — between organizations that governed their acceleration and those that merely survived it — is going to matter more than most people currently believe. And by the time the market makes that unmistakably clear, the window to establish it will have already narrowed.

About Kumiho Strategies

Kumiho Strategies is a boutique federal advisory and delivery firm specializing in operational governance, acquisition support, modernization strategy, and AI-assisted knowledge work. This paper is part of the Kumiho Operational Perspectives series — analysis written from inside the federal modernization environment, not from outside it.

kumihostrategies.com

ENDNOTES

Sources

[1] NOAA, "NOAA Deploys New Generation of AI-Driven Global Weather Models," December 17, 2025. noaa.gov/news-release/noaa-deploys-new-generation-of-ai-driven-global-weather-models. NOAA is an agency within the U.S. Department of Commerce. Corroborated by HPCwire, TechSpot, and CBS News.

[2] Ono, H. et al., "Autonomous Robotics Is Driving Perseverance Rover's Progress on Mars," Science Robotics (peer-reviewed): AutoNav used for 88% of first Mars year travel. IEEE Spectrum, February 2026: more than 90% autonomous travel confirmed as of October 2024. NASA JPL, "NASA's Perseverance Rover Completes First AI-Planned Drive on Mars," January 30, 2026. jpl.nasa.gov

[3] State Magazine, U.S. Department of State, "AI in Action," December 2024. statemag.state.gov. StateChat launched August 2024; 1,900+ first-time users and 10,000+ prompts in first week. state.gov/artificial-intelligence lists StateChat as deployed enterprise-wide for SBU content. FedScoop, "State Department is gearing up to roll out agentic AI," February 13, 2026.

[4] GAO-25-107653, Generative AI Use and Management at Federal Agencies, July 29, 2025. gao.gov/products/gao-25-107653. Brookings Institution analysis, April 2026: 3,600+ use cases, 69% increase over 2024, approximately five times 2023 levels. FedScoop, "Disclosed Government AI Use Increased by 70% in 2025."

[5] GSA, "Launching the OneGov Strategy, for IT," April 30, 2025. gsablogs.gsa.gov. Washington Technology, "GSA hits one-year mark for OneGov," April 2026: $1.1 billion in savings, 20 agreements. FedScoop, "GSA's OneGov strategy won quick raves from federal IT leaders. Will it have staying power?," February 25, 2026: "praise, but also confusion over how the program works."

[6] Stats NZ, Regulatory Impact Statement: Modernising the Census, 2025. stats.govt.nz. NZ Council for Civil Liberties, Submission: Future Census, June 2024; Submission: Data and Statistics (Census) Amendment Bill, April 2026. nzccl.org.nz. The Conversation, "Data gaps and demographic change: the end of the NZ census will create big blind spots," January 20, 2026. Statutory Review of New Zealand's 2023 Census, Newsroom, April 29, 2024.

[7] OMB 2024 Federal AI Use Case Inventory, primary data. github.com/ombegov/2024-Federal-AI-Use-Case-Inventory (file: 2024_consolidated_ai_inventory_raw_v2.xls). Direct analysis: 2,133 total; 351 flagged rights/safety-impacting; 125 extensions in raw data. OMB December 2024 published release cited 206 extensions granted. BABL AI analysis, September 2025: confirmed 206 extensions, noting they do not auto-renew. VA: 145 rights/safety systems, 109 extensions. SSA: 9 rights/safety systems, all public-facing HISP services, all 9 requested extensions.

[8] KrebsOnSecurity, "CISA Admin Leaked AWS GovCloud Keys on GitHub," May 2026. krebsonsecurity.com. Repository created November 13, 2025; accessible until May 15, 2026 (~6 months). Contractor employed by Nightwing, Dulles VA. Discovered by Guillaume Valadon, GitGuardian. Corroborated by The Stack, Cybernews, CyberPress, and axis-intelligence.com technical analysis.

[9] Brookings Institution, Federal AI Use Case Inventory Analysis, April 2026: "More than 85% of high-impact deployed AI use cases in 2025 lacked some required information." fedscoop.com/disclosed-government-ai-use-increased-in-2025-omb; executivegov.com/articles/omb-federal-ai-use-cases-2025.