

SUN RADIA

Distributed AI for the Mission

Operationalizing AI in the Public Sector
and Deploying It at the Power-Fiber Edge

*From standard operating procedures to scalable
agentic workflows —and the sovereign
infrastructure to run them at mission speed.*



business
of data



Contents

Click below to navigate

INTRODUCTION

Overcoming Public Sector Pilot Paralysis

1

CHAPTER ONE

Operationalizing AI in the Public Sector

2

CHAPTER TWO

*Edge AI Factory:
Distributed AI Infrastructure for the Mission*

10

CHAPTER THREE

Bringing It to Award: The Tradewinds Path

13

YOUR NEXT STEP

AI Modernization from Pilot to Production

16

Contributors



Gabe Arrington
Chief, International
Affairs, United States
Air Force



Steven Walters
CEO, Property
Management
Resources LLC



Richard Patterson
Chief Data and AI
Officer, Bureau of the
Comptroller and Global
Financial Services,
U.S. Department of State



INTRODUCTION

Overcoming Public Sector Pilot Paralysis

For federal agencies, the central challenge of artificial intelligence is no longer whether a model can work — it is how to operationalize it. Promising pilots routinely stall on the way to production because they must integrate with legacy systems, run on messy and tightly governed data, satisfy security and audit requirements, and earn the trust of the mission users who depend on them. The problem is compounded by infrastructure: most AI is built to run in centralized, hyperscale data centers that sit far from where federal missions operate, introducing latency, sovereignty, and resilience gaps that make real-time, edge-driven decisions impractical.

Addressing this means moving both the data and the AI to the edge. A distributed infrastructure model places containerized compute where power and fiber already converge — close to where the mission actually runs — so that

data is processed and inference is served on-site rather than shipped to a distant cloud. Keeping data and AI at the edge delivers sub-10ms performance, federal-grade sovereignty, and geographic resilience, directly removing the operational barriers that leave AI stranded in the pilot phase.

This paper brings together three companion pieces that trace that arc from challenge to solution. The first, drawn from conversations with public sector data and mission leaders, examines what it takes to turn standard operating procedures into scalable, governed agentic workflows. The second and third describe Edge AI Factory — a distributed AI infrastructure platform purpose-built to run those workloads at the mission edge — and the path to fielding it through the Tradewinds Solutions Marketplace. Each chapter preserves its source material and can be read on its own. ■



CHAPTER ONE

1

Operationalizing AI in the Public Sector

*Turning Standard Operating Procedures
into Scalable Agentic Workflows*

CONTRIBUTORS

Gabe Arrington

Chief, International Affairs, United States Air Force

Richard Patterson

Chief Data and AI Officer, Bureau of the Comptroller
and Global Financial Services, U.S. Department of State

Steven Walters

CEO, Property Management Resources LLC

The Operationalization Challenge

For public sector leaders, AI rarely arrives in a pure environment.

Often, organizations' data infrastructure is shaped by legacy systems, tight budgets, and strict approval chains, and every modernization decision must be made with public trust in mind.

None of this can be paused while organizations experiment. Leaders must innovate, but they must also protect trust. They must move faster, but they must preserve accountability.

They must encourage experimentation, but guard public services.

That makes the challenge of operationalizing AI fundamentally different from proving that AI can work in a controlled setting. A pilot can show that a model can perform discrete tasks or provide information. Production AI must do far more. It must integrate with existing systems, operate on messy data, fit within governance frameworks, support auditability, and earn the confidence of the people expected to use it. ▶



This is why many public sector AI initiatives stall between proof of concept and operational value. The challenge, then, is not whether we can make technology work, it is how to best prepare our organizations to benefit from it.

For Gabe Arrington, Chief, International Affairs for the United States Air Force, the distinction between an interesting AI experiment and an operationally useful capability boils down to three fundamentals: integration, sustainment, and mission alignment.

“The barrier most frequently underestimated is the sustainment tail, specifically machine learning operations and continuous monitoring. Leaders often treat AI like traditional commercial-off-the-shelf technology or software. You buy it, you install it, you move on. In reality, AI models degrade over time as real-world data shifts from training data.”

Gabe Arrington

Chief, International Affairs,
United States Air Force

That point is especially important in public sector environments, where AI systems must operate within established controls rather than around them. A technically successful prototype can still fail

to create value at scale if it cannot be appropriately monitored, governed, secured, and sustained as conditions change.

Richard Patterson, Chief Data and AI Officer, Bureau of the Comptroller and Global Financial Services for the U.S. Department of State, argues that many AI initiatives struggle because leaders underestimate the full spectrum of support that takes AI from experiment to production scale.

“Someone has a great idea and someone who has some authority likes that idea and they say, let’s go do it. Something gets built and it may not be exactly what the end user needed. It may not fit into the enterprise architecture. It may not have security and governance around the data that needs to be there.”

Richard Patterson

Chief Data and AI Officer,
Bureau of the Comptroller
and Global Financial Services,
U.S. Department of State

The result is a familiar pattern. A team builds something that demonstrates technical promise, but the tool does not align with how work actually gets done. It may depend on data that is inconsistent, poorly governed, or difficult to access. ►



It may require users to change long-established workflows without giving them enough transparency or incentive to trust the output. It may solve one step in a process while creating friction elsewhere.

In the public sector, these issues are made more pronounced by organizational complexity. Processes often span multiple teams, systems, and approval points. Responsibilities may be distributed across departments.

This means public sector organizations cannot treat AI as a shortcut around foundational work. Scaling AI requires a clear understanding of the mission outcome, the data environment, the enterprise architecture, the decision rights, the risk profile, and the human oversight model.

It also requires a budget and talent strategy for sustainment, because AI systems are not static assets that can be installed and left alone. AI should not be inserted into isolated process steps without considering the wider workflow.

The organizations that move beyond AI pilots reliably will be the ones that connect experimentation to mission outcomes, architecture, governance, workflow design, funding, and change management from the start.

This is the foundation for operational AI. Before agencies can turn standard operating procedures into scalable agentic workflows, they must first close the gap between what AI can do in theory and what public sector organizations need it to do in practice. ►

Are Your Standard Operating Procedures Ready for Agentic Workflows?

Beyond scaling promising AI use cases, public sector data leaders need to identify workflows that can turn repeatable work into governed workflows that can operate across systems, teams, and approval chains without losing accountability.

That is where standard operating procedures, or SOPs, become critical. In theory, SOPs show leaders where AI can assist, where people must stay in control, and how decisions should move through an organization. In practice, many public sector organizations first have to confront an uncomfortable reality: not every process is as standardized as leaders assume.

According to Richard Patterson, many teams discover this only when preparing a process for automation.

“A lot of the business types of activities, surprisingly, are done through word-of-mouth processes. There might be some training materials, there might be an SOP, it might be 20 years old.”

Richard Patterson

Chief Data and AI Officer,
Bureau of the Comptroller
and Global Financial Services,
U.S. Department of State

This matters because agentic workflows depend on clarity. If a process varies across teams, or exists mainly as institutional knowledge, it becomes difficult to automate responsibly.

Before leaders can decide whether AI should summarize a document, populate a form, trigger an authorization, or escalate a decision, they must understand the real workflow, not just the assumed one.

For Patterson, having a documented SOP is a starting point. From there, leaders can evaluate risk, reward, governance needs, and guardrails. A process with high volume, repeatable steps, and clear escalation points may be a good candidate for AI-enabled orchestration. A process with unclear ownership, inconsistent execution, or high-impact discretionary decisions may need redesign before automation. ►

Steven Walters, CEO of Property Management Resources LLC, adds that readiness also depends on how an organization defines success up front.

“A workflow is ready for agentic AI when its decision logic is documented, its inputs are clean and governed, and the work has a clear definition of done, an auditable success criterion and a defined exception path. The strongest warning sign is when a process is tribal, executed a little differently by every person. Automating that only scales the chaos. The deepest red flag is intent: using automation as a band-aid to avoid fixing a broken process.”

Steven Walters

CEO, Property Management Resources LLC

Gabe Arrington says leaders should think of AI as a way to orchestrate established procedures, not bypass them. That distinction is especially important in government environments, where approvals, documentation, handoffs, and records often exist for good reasons.

“The goal is to use AI to orchestrate the standard operating procedure, not to bypass it. Accountability is maintained by shifting the human from doer to reviewer. You establish control through robust auditability, every AI action, data retrieval, suggested decision must be logged.”

Gabe Arrington

Chief, International Affairs,
United States Air Force

This is the practical promise of agentic AI in public sector environments. Rather than replacing people, agentic systems can coordinate work across fragmented systems and teams.

In doing so, they can reduce the manual “swivel chair” work that slows many government processes, while keeping sensitive decisions under human authority. Arrington describes agentic AI as connective tissue between legacy systems that were never designed to communicate.

This makes workflow orchestration especially relevant for compliance, authorization, data cataloging, and pipeline management. These are areas where work is often structured, documentation-heavy, and dependent on multiple systems. ►



But orchestration is not just a technical layer. It is also a control mechanism. By designing workflows around audit trails, permission checks, escalation points, and human review, agencies can move faster without weakening governance. This is essential in public sector environments, where speed alone is not enough. AI-enabled workflows must also be explainable, traceable, and legally defensible.

The human role also changes. In an orchestrated workflow, AI can absorb more of the administrative burden, while people focus on higher value tasks.

However, the shift only works when users trust the system. Patterson warns that if people do not understand where AI-generated information comes from or how AI-influenced decisions are made, it may damage that trust.

For Steven Walters, CEO of Property Management Resources LLC, the way to keep human authority meaningful without slowing the process down is to make oversight proportional to risk.

“The way to keep authority meaningful without creating bottlenecks is to make oversight risk-based rather than uniform. Low-risk, repeatable steps run autonomously within clear guardrails, while human review is reserved for consequential judgments and anything that crosses a defined threshold. That’s what turns oversight into a control mechanism embedded in the workflow, not a manual gate bolted onto it.”

Steven Walters

CEO, Property Management Resources LLC

This is why successful AI orchestration must involve end users early. Their job descriptions and performance metrics may also need to evolve so they are not effectively competing with AI for the same task.

Steven Walters, CEO of Property Management Resources LLC, argues that this involvement has to start earlier than most organizations plan for.

“Leaders should involve frontline employees early, during process discovery, not after the workflow is designed. The people doing the work know how a process actually runs versus how it’s assumed to run. That’s also how you address the fear directly: be explicit that the goal is to shift people from doers to reviewers, not to compete with them for their own tasks.”

Steven Walters

CEO, Property Management Resources LLC

Turning SOPs into agentic workflows is therefore not a simple automation exercise. It is a modernization discipline that combines process discovery, governance, enterprise architecture, data management, and change leadership.

This offers a tantalizing opportunity to increase efficiency in the public sector. But the sequence matters. Leaders must first understand the process, then design the workflow, then apply AI within clear boundaries.

Only then can agentic AI help public sector organizations scale operational improvements while preserving public trust. ►



How to Extend Orchestration Across Modernization Initiatives

Once public sector organizations have the necessary foundations for growth, workflow orchestration can support the scaling of automation across the organization in a consistent way.

Many modernization initiatives are not limited by a lack of AI capability. Rather, it's the fragmented systems, manual handoffs, inconsistent data flows, and processes that lie underneath that slow them down. In these environments, AI may provide insight, but workflow orchestration turns that insight into structured action.

Arrington argues that workflow orchestration is the layer that makes AI operationally useful. Without it, even powerful algorithms remain trapped in isolated systems.

“Workflow orchestration is really the foundational layer that makes AI usable. Without it, you simply have brilliant algorithms created by the greatest minds in the world trapped in isolation.”

Gabe Arrington

Chief, International Affairs,
United States Air Force

This is especially important for processes that depend on traceability and repeatable execution. They also require agencies to gather evidence from multiple systems, apply rules, document decisions, and escalate exceptions to the right people.

Patterson expects this model to become increasingly important as AI products become more domain-specific. In areas such as government finance, logistics, compliance, and back-office operations, he sees opportunities for AI-enabled platforms that incorporate common processes and best practices.

“There’s going to be more domain-based products as well, that understand the best practices and the optimal processes to get the most efficient ways out.”

Richard Patterson

Chief Data and AI Officer,
Bureau of the Comptroller
and Global Financial Services,
U.S. Department of State

But the path forward is not simply to buy or build more AI tools. Public sector leaders need to decide where standardized products are appropriate, where workflows should be adapted, and where mission-critical capabilities require custom development. This is where enterprise architecture becomes central to modernization.

Ad hoc experimentation is not enough. If agencies want AI to scale, they must connect it to their future-state architecture, data strategy, governance model, and investment plan. ►



“Let’s just throw AI spaghetti at the wall and see what sticks, is not a strategy. Incorporating AI into your architecture is a strategy, and that’s where you develop your governance and your spend.”

Richard Patterson

Chief Data and AI Officer,
Bureau of the Comptroller
and Global Financial Services,
U.S. Department of State

This strategic view is essential because orchestration can extend across many modernization priorities at once. A framework developed for compliance evidence gathering may also support ATO workflows. A data cataloging workflow may improve pipeline governance. A process for routing AI-generated outputs for human review may become the basis for broader responsible AI controls.

The opportunity is to create reusable patterns for modernization. Instead of treating each AI use case as a separate project, agencies can build common orchestration capabilities that support multiple functions. Over time, this can help reduce duplication, improve consistency, and make governance easier to enforce.

Steven Walters, CEO of Property Management Resources LLC, frames this as a shift to platform-oriented approach.

“Organizations avoid rebuilding governance for every use case by treating AI as a platform capability rather than a series of one-off projects. Audit trails, permission checks, escalation points, and human review become standard building blocks that any new workflow inherits by default, so only the mission-specific logic changes from case to case.”

Steven Walters

CEO, Property Management
Resources LLC

The clearest sign of progress is not the number of pilots launched, or AI tools procured, but whether those capabilities are becoming part of everyday mission delivery. It also changes how leaders should measure the maturity of AI integration.

“We’ll know we’ve moved beyond the hype when we stop talking about AI initiatives and simply talk about mission outcomes.”

Richard Patterson

Chief Data and AI Officer,
Bureau of the Comptroller
and Global Financial Services,
U.S. Department of State

That is the long-term promise of workflow orchestration in the public sector. Making AI part of the operating fabric of government in support of faster, safer, and more accountable modernization. ■



2

CHAPTER TWO

Edge AI Factory:

Distributed AI Infrastructure for the Mission

*Distributed AI infrastructure for the
mission, at the power-fiber edge*

Tradewinds Solutions Marketplace · AI / ML · Federal Infrastructure · May 2026

Executive Summary

Edge AI Factory is a distributed AI infrastructure platform that deploys containerized compute at the seam of utility substations and telecom central offices. It delivers sub-10ms inference, federal-grade resilience, and domestic sovereignty — at the speed the mission requires, not the speed hyperscale construction queues allow. ►

The Problem

Power, not capital, is now the binding constraint on AI infrastructure. More than 2,000 GW of generation and storage — nearly twice the nation's installed capacity — sit stalled in U.S. interconnection queues, and the median project now waits roughly five years from request to operation, up from under two years in 2008 (Lawrence Berkeley National Laboratory). In the PJM territory that serves the National Capital Region, projects reaching operation in 2025 had waited an average of eight years.

Against that backlog, AI data-center demand is projected to grow 3.5x between 2025 and 2030 to roughly 156 GW worldwide (McKinsey), and a single AI data center can draw 200+ MW versus 30–50 MW for a traditional facility — producing a 3–5 year construction-and-interconnection timeline the mission cannot wait out.

Centralization compounds the problem. Hyperscale campuses sit hundreds of miles from where federal missions operate, imposing greater than 50 ms of round-trip latency — untenable for workloads where decisions are measured in milliseconds — while concentrating risk into a small number of large, remote footprints.

Inference, not training, now drives these economics: it accounts for 80–90% of the lifetime cost of a production AI system because it runs continuously, and every query shipped to a distant cloud adds latency, egress cost, and sovereignty exposure.

Department of War (DoW) workloads — ISR fusion, C2 decision support, autonomous edge AI — reward proximity, sovereignty, and modularity over raw scale.

The Solution

Each AI Factory is a containerized deployment unit installed at sites where electrical substations and telecom fiber already converge.

A four-tier architecture — Cloud, 5G/MEC, Far Edge, and On-Device — distributes inference where it is needed, with phased utility interconnection that aligns capital with contracted demand.

The flagship site at Fort Washington, Maryland is owned, zoned, approximately 0.4 miles from a Pepco substation, adjacent to a Verizon Central Office, and inside the Joint Base Andrews / National Capital Region (NCR) corridor. ▶

Mission Outcomes Rolling Q2

99.6%

Latency-met
— sub-10ms
inference SLA

8.4m

Real-time
inference calls
accelerated

3 / 12

Sites operational
— live · contracted
pipeline

Why Now & Path Forward

Three trends now converge to make distributed, edge-based AI the right model at the right time. First, inference has become the dominant workload: it already represents roughly two-thirds of all AI compute in 2026 — up from one-third in 2023 and half in 2025 — and Gartner projects 55% of AI-optimized IaaS spending will support inference this year, rising above 65% by 2029 (Deloitte; Gartner).

The AI inference market is on track to grow from about \$106 billion in 2025 to \$255 billion by 2030 (roughly 19% CAGR), and the edge-AI segment specifically from \$24.9 billion to \$66.5 billion over the same period (21.7% CAGR). Inference rewards proximity, security, and reliability over raw scale — exactly what a distributed footprint delivers.

Second, power has replaced capital as the binding constraint, and

utilities increasingly favor smaller, phased interconnections at existing substations over multi-gigawatt greenfield hookups mired in five-to-eight-year queues.

Third, federal policy pulls in the same direction: DoW, CDAO, and federal AI mandates explicitly require domestic, zero-trust, geographically resilient compute, and the Department of Energy's October 2025 directive asserts federal jurisdiction over new large loads above 20 MW to standardize and accelerate their interconnection.

Edge AI Factory is post-competition and awardable through the Tradewinds Solutions Marketplace, capable of going from task order to a live AI Factory in under 120 days — against the 3–5 year hyperscale timeline. The platform is ready to scale across the National Capital Region, the Mississippi (TVA) corridor, and additional federal-priority corridors. ■

**Time-to-mission, not time-to-build.
Bring it to mission. Bring it to award.**

3

CHAPTER THREE

Bringing It to Award: The Tradewinds Path

Federal Tradewinds Edition · Tradewinds Solutions Marketplace · May 2026

Post-Competition and Awardable Today

Edge AI Factory is available to federal mission owners now through the Tradewinds Solutions Marketplace. It is post-competition and awardable, which means government customers can move from requirement to task order without running a new full-and-open competition. The platform is engineered to go from task order to a live AI Factory in under 120 days — replacing the 3–5 year hyperscale construction timeline with a time-to-mission model.

The flagship site in Fort Washington, Maryland is owned, zoned, approximately 0.4 miles from a Pepco substation, adjacent to a Verizon Central Office, and within the Joint Base Andrews / National Capital Region (NCR) operational corridor — a repeatable template for siting at power-fiber convergence points inside federal-priority corridors. ▶



SUN RADIA

[Talk to Sun Radia about your mission](#)

Why Now — Three Converging Tailwinds

1

Inference is becoming the dominant workload

By 2027, inference is projected at 60–70% of AI compute spending — rewarding proximity, security, and reliability over raw scale.

2

Power has replaced capital as the binding constraint

Major U.S. data center markets face 3–5 year power interconnection backlogs; distributed deployment at existing substations is the only scalable response.

3

Federal mandates favor distributed, sovereign compute

DoW, CDAO, and federal AI requirements explicitly demand domestic, zero-trust, geographically resilient infrastructure — procurable now through Tradewinds.

The Platform

A four-tier architecture (Cloud, 5G/MEC, Far Edge, On-Device) is deployed in containerized modules at power-fiber convergence sites. Phased utility interconnection (P1 5 MW → P2 15 MW → P3 30 MW → P4 60 MW) aligns delivered capacity with contracted mission demand.

Each site is a repeatable, vertically integrated template: substation adjacency, telco-grade fiber, and federal-corridor positioning — giving mission owners a sovereign, resilient footprint rather than a share of a distant hyperscale campus. ►

Mission Outcomes Rolling Q2

99.6%

Latency-met
— sub-10ms
inference SLA

8.4m

Real-time
inference calls
accelerated

3 / 12

Sites operational
— live · contracted
pipeline

The Acquisition Advantage

For federal buyers, the value of the Tradewinds path is speed and fit. Because the solution is already competed and available in the marketplace, mission owners avoid lengthy source-selection cycles and can direct capability to where it is needed.

The distributed, containerized model means capacity can be added corridor by corridor as requirements grow, with each AI Factory delivered as a self-contained, zero-trust unit that satisfies domestic-sovereignty and geographic-resilience requirements out of the box.

Path Forward

Edge AI Factory is ready to scale across the National Capital Region, the Mississippi (TVA) corridor, and additional federal-priority corridors.

The next step for any mission owner is a task order: from award, a live AI Factory can be standing up inference at the edge in under 120 days. ■

**Time-to-mission, not time-to-build.
Bring it to mission. Bring it to award.**

YOUR NEXT STEP

AI Modernization from Pilot to Production

Operationalize AI—and deploy it at the power-fiber edge, Fed Ready Today.

Federal agencies run support processes on standard operating procedures that are ready for agentic automation, yet agency initiatives face two barriers to production: meeting federal compliance requirements and running workflows where the work happens and the data lives, not in a centralized location.

Sun Radia closes this gap with a production-ready agentic framework that turns SOPs into governed agentic workflows, running on sovereign edge-to-cloud infrastructure built for the mission.





What We Bring

Operationalize AI

We turn your SOPs into scalable, governed agentic workflows—orchestrating work across legacy systems with audit trails, permission checks, and human-in-the-loop review at every consequential step.

Edge AI Factory

Containerized, sovereign compute deployed where power and fiber already converge. Sub-10ms inference, federal-grade resilience, and domestic data sovereignty—at mission speed.

Fed Ready Now

Assessed and awardable through the Tradewinds Solutions Marketplace—move from requirement to task order without running a new full-and-open competition.

Inference is the Dominant workload

60–70% of AI compute by 2027
—rewarding proximity, security, and reliability over raw scale.

Power is the New Capital

3–5 year interconnection backlogs, make distributed deployment at existing substations the scalable response.

Federal Mandates Favor Sovereign Compute

DoW, CDAOs, and Federal AI rules require domestic, zero-trust, geographically resilient infrastructure.

99.6%

Latency-met
— sub-10ms
inference SLA

8.4m

Real-time
inference calls
accelerated

<120

Days from task order
to a live AI Factory

4 weeks

Scope to
Production

Start Here - Three Low-risk Steps

1

Scope a use case

A one-week diagnostic to identify a high-value, automation-ready workflow—and confirm the data and decision rights are in place before we build.

2

Prove it in one week

A governed proof of concept on your data, in your environment—guardrails and human oversight built in from day one.

3

Production in two weeks

Move to a live AI Factory and add use cases on the same governed stack as your mission grows—no re-platforming.

TALK TO US

Robert Keenan

Robert.Keenan@SunRadia.com
703-927-9787

Sudha Gummuluru

Sudha.Gummuluru@SunRadia.com
240-997-5996

www.SunRadia.com

via Tradewinds Solutions
Marketplace GSA MAS
engage@SunRadia.com



Discover Corinium Intelligence

Corinium is a specialist market intelligence and events company, with the world's largest business community of more than 300,000 data, analytics, customer experience and digital transformation leaders.

Through insights, practical tools, expert guidance and industry connections, we help executives drive efficiencies and innovation.

Find out more: www.coriniumintelligence.com

Connect with Corinium

 [Join us at our events](#)

 [Visit our blog](#)

 [Read our white papers](#)

 [Follow us on LinkedIn](#)