WHY **UNIFIED DATA** IS THE FOUNDATION **FOR AI-READINESS**

Agencies moving beyond 'AI 1.0' are adopting the enhanced performance and lower costs of a unified data management platform, delivered as a service.

By Scoop News Group

Federal agencies are racing to harness the speed and capabilities of artificial intelligence. The latest figures from OMB indicate more than 3,600 AI use cases are underway across the government — a 70% increase over the previous year. That doesn't include projects in progress at the Department of Defense, most notably the department's recently announced $30 billion "AI Arsenal initiative."

However, behind the rapid rise in use cases and pilot projects lies a sobering reality.

Getting AI to deliver meaningful enterprise value is proving brutal, not just in government, but across multiple industries.

- A RAND Corporation report estimated that 80% of AI projects fail due to a combination of factors — a failure rate twice that of IT projects not involving AI. The root causes include poor data quality, underinvestment in infrastructure, weak risk controls, and a failure to demonstrate return on investment.
- A Gartner survey similarly found that at least half of all generative AI initiatives had collapsed after the proof-of-concept phase through last year. A key reason: difficulties overcoming entrenched data management practices. In another report, Gartner predicted that through 2026, organizations will abandon 60% of AI projects that lack AI-ready data.

The chief bottleneck isn't the AI models, say experts like Dan Kent, public sector CTO at Everpure — it's primarily a broken data foundation.

"The government's deploying AI everywhere," says Kent, in what he refers to as "the 1.0 release of AI. Agencies are putting it out there but not necessarily getting it right." Moving beyond "AI 1.0" will require agencies to gain "visibility into where their data originated, how it was transformed, whether it remains authorized for use and how it is being applied to AI systems," Kent said in a recent op-ed.

KEY CHOKE POINTS

Why is AI 1.0 stalling, especially at government agencies? Kent sees three primary culprits:

- **Misguided alignment.** Many AI pilot projects are often launched to meet top management's requirements without fully factoring in the complexity and shortcomings of the agency's broader IT ecosystem and mission demands.
- **Fractured data foundation.** Government data remains trapped in complex, legacy silos. When an agency wants to complete a mission task

or outcome, it typically must pull or copy data from multiple sources with conflicting ownership and governance rules. That's a recipe for massive data reliability problems.

- **Pilot purgatory:** Even when agencies have modern data management principles in place, they invariably hit a financial wall when trying to scale. Between the eye-watering costs of the GPUs required to build models — or the token allowances to run AI queries — moving from a conceptual pilot to a daily operational application triggers massive sticker shock. Even if funding is available, the talent to integrate AI into existing operations is often not.

Kent points to two other challenges:

- **Security vulnerabilities:** Copying and consolidating data from different sources to an intermediary cloud or on-premises environment for AI inferencing introduces a host of recurring security and privacy issues.
- **Self-inflicted hallucinations:** In a dynamic environment, copied data goes stale instantly. "By making copies of data to put into your AI platform, you now have a built-in method of making that data outdated and not in sync with the raw data, which can result in self-inflicted hallucinations or misleading outcomes," Kent explains.

FIXING THE DATA FOUNDATION

While the rise of AI promises unprecedented operational and mission breakthroughs, it has also unleashed an exponential wave of data that traditional storage systems are hard-pressed to manage. Instead of scaling up, legacy infrastructure is breaking down — leaving organizations trapped in a maze of fragmented silos and data sprawl.

> **"By using a unified data plane with a software-defined operating environment, agencies can prepare data for AI at the source without adding more fragmented infrastructure or management silos."**
>
> – Dan Kent, Public Sector CTO, Everpure

The need to rethink data storage infrastructure isn't just about addressing the growing volume of data; it's also about how AI is disrupting data management itself.

"Legacy databases are becoming a smaller fish in a bigger AI pond … and being recast as AI databases," observes John Foley, a database analyst and founding editor of the Cloud Database Report. "Data architectures are increasingly decentralized, with data mesh, fabrics, data clouds, multi-cloud databases, edge devices, and more. Data sovereignty and resiliency requirements are forcing new thinking around how and where data gets stored."

"There's something else at play," he adds. "More databases are being created autonomously by AI agents. As this happens, it points to a net increase in the sheer number of databases, driven by agentic AI." All of which helps explain the architectural shift underway.

EMBRACING UNIFORM, INTELLIGENT DATA PLANES

That's a key reason why organizations, including NASA and the Department of Defense, are moving to a different approach to data management and storage and increasingly recognizing the value of a unified, centrally managed enterprise data platform delivered as a service.

"AI-ready data is not 'one-and-done,'" argues Roxane Edjlali, senior director analyst at Gartner. "Think of it as a practice where the data management infrastructure needs constant improvement based on existing and upcoming AI use cases."

That requires moving beyond thinking of storage as a passive utility and embracing an intelligent, unified data plane that provides full visibility, access and control across the entire data estate — on-premises, in the cloud and at the edge, she says.

Gartner predicts that by 2029, consumption-based storage as a service (STaaS) will replace 50% of on-premises enterprise storage and data services infrastructure capital expenditure (capex), an increase from 15% in early 2025.

WHY IT MATTERS

Unlike current market practices, a unified data plane supported through an enterprise storage cloud as a service directly addresses the primary operational pain points of most federal IT leaders, including:

- **Crushing the power and space crisis**
  Modern AI workloads are consuming data center resources at an alarming rate. A unified, flash-based architecture completely rewrites the real estate equation. "Most of our customers... struggle just for power and space," Kent notes. "Within one rack, we can host a lot of the capabilities that would take a dozen or two dozen racks by our competition."
- **Ending "forklift upgrades"**
  Traditional data models trap agencies in endless three-to five-year overprovisioning cycles. A consumption-aligned service model can help agencies scale more incrementally, improve budget predictability, and reduce the operational disruption associated with major refresh cycles.
- **Reducing data risks**
  "By using a unified data plane with a software-defined operating environment, agencies can prepare data for AI at the source without adding more fragmented infrastructure or management silos," Kent explains. "It also avoids replicating and moving data that can create 'data drift' scenarios as well as introduce hidden security risks."
- **Mitigating the skills shortage**
  Federal IT shops are notoriously understaffed. A unified data plane standardizes a complex, multi-vendor environment, running data tiering and continuous optimization automatically in the background. The result? A consistent, cloud-like data experience where complex security workflows are deployed at the "push of a button."

ACCELERATING YOUR DATA READINESS

To build an infrastructure and data foundation capable of scaling AI beyond a pilot, Kent suggests that command and director-level leaders should focus on five key strategies:

1. **Consolidate the architecture:** Unify structured database systems with unstructured file and object data. Apply a single policy system across the entire fleet to eliminate the operational complexity of multi-vendor environments.
2. **Automate the embedding pipeline:** Use software-defined data streams to tokenize and index raw records into vector embeddings automatically. Ensure that any update to the source data triggers an automatic refresh of the vector database to eliminate model inaccuracies.
3. **Build an "AI Factory":** For highly sensitive or classified workloads, deploy high-throughput, on-premises infrastructure. Your storage must deliver multiple terabits per second to properly feed GPUs during large-scale model training and inferencing.
4. **Enforce source-level governance:** Don't wait for the application layer to secure data. Implement tools that provide deep discovery and semantic context—identifying Personally Identifiable Information (PII) and sensitive records directly at the source before the data is fed to the AI engine.
5. **Shift the Economic Model:** Move away from massive, rigid Capital Expense (CapEx) hardware buys that lead to over-provisioning. Adopt flexible, utility-based Operational Expense (OpEx) subscription tiers to map technology spending directly to data consumption.

SUPPLY-SIDE EVOLUTION

Everpure's own evolution exemplifies the industry's shift toward data management. Formerly known as Pure Storage, the company has spent the past five years pivoting from core storage arrays to advanced data management. It was an early mover in developing service-led storage models and tools to unify data, and its efforts have gained increasing attention.

Gartner recognized Everpure as a top "Leader" for execution and vision this past year in its Magic Quadrant for Enterprise Storage Platforms.

In particular, Gartner cited Everpure's Evergreen//One platform, built on FlashArray and featuring an intelligent control plane for automation and self-optimizing performance, for delivering high operational efficiency, unified data management and strong customer satisfaction.

"By tightly integrating software with Everpure's DirectFlash architecture over raw NAND flash, rather than relying on legacy disk-based storage stacks, Everpure can deliver the throughput needed to keep large AI training environments fed while reducing footprint, power, and overall cost." Kent noted. "Everpure has roughly five

> **Legacy databases are becoming a smaller fish in a bigger AI pond. Data sovereignty and resilience requirements are forcing new thinking around how and where data gets stored."**
>
> - John Foley, database analyst

times the capacity that most larger data centers have, in half the footprint," he said. "Moreover, we're the third largest acquirer of NAND flash memory, which puts us in a good place, given the issues around memory nowadays.

Crucially, Everpure has tightly integrated its data streaming software with the **Nvidia Blackwell GPU** architecture. This joint partnership allows agencies to build high-performance "AI Factories": on-premises, creating an automated, data-ready pipeline that scales inferencing safely across the enterprise," Kent said.

To strengthen its position in the AI ecosystem, Everpure recently **acquired 1touch**, a specialized technology that delivers advanced data discovery and contextual semantic mapping. This capability allows agencies to scrub data and identify exactly what sensitive information exists before it ever reaches an LLM.

STREAMLINING THE DATA EXPERIENCE

What also distinguishes Everpure's position with federal customers, according to Francesca El Attrash-Ukaejiofo, senior product marketing manager for public sector at Everpure, is its focus on delivering a "seamless data experience."

"Everpure delivers an easy-to-use data experience across all environments, whether it's managing data in the cloud, on premises, or at the edge — it's one unique, consistent experience," she said.

The clear lesson of AI 1.0, concludes Kent, is that an AI application is only as good as the data feeding it. Pursuing advanced algorithms while leaving data rotting in fractured, stagnant silos is an expensive exercise in futility.

The mandate for federal leadership is straightforward: "Prioritize data readiness over model deployment. Stop funding isolated AI tools that trap the agency in pilot purgatory. Shift your strategic focus to building an intelligent, unified data plane. Fix the foundation first, and ensure that the autonomous systems guiding tomorrow's civilian and defense missions are drawing from an unassailable source of truth.

Learn how Everpure is helping government agencies and enterprises lay the right data foundation for the AI era.

This report was produced by Scoop News Group and sponsored by Everpure.

Everpure

SCOOP NEWS GROUP