Everforth™
ECS Federal®

2026 FEDERAL DATA AND AI REPORT

Foundations, Frontiers, and Fault Lines

A 12-Month Look Forward Into the Federal Data and AI Landscape

By Matthew McDonald, Austin Amaya, Ph.D., Ketan Mane, Ph.D., Mojgan Pedoeim, Patrick Elder, Anthony Zech, and Harkiran Kaur

VOL 2026 | Issue No. 2

Table of Contents

INTRODUCTION 3

FOUNDATIONS 5

1. Evolving AI Ecosystems: From Tools to Orchestrated Systems 5
2. Data Readiness Still Determines Success 6
3. AI Governance, Compliance, and the New Procurement Reality 7

FRONTIERS 8

1. Mission-First AI: The Rise of Vertical Intelligence 8
2. Agentic AI: From Assistants to Autonomous Systems 9
3. Speed to Impact vs. Trust and Validation 10
4. Enabling the AI-augmented Workforce 11

FAULT LINES 12

1. Strategy Gaps and the Persistence of Pilot Purgatory 12
2. Compute, Cost, and Infrastructure Constraints 13
3. The AI Reality Gap: Navigating Slop vs. Value 14
4. Trust, Security, and Explainability Risks 15
5. The AI Talent and Literacy Gap 16

CONCLUSION 17

Introduction

This report provides a practical, forward-looking view of the federal data and AI landscape over the next 12 months. Structured around **foundational capabilities, evolving priorities, and the fault lines that can undermine success**, *it outlines where agencies should focus now to maximize the value of AI investments and prepare for continued transformation.*

Organizations Want to Move Beyond AI Pilots. Scaling Challenges Remain.

Artificial intelligence has moved beyond experimentation in many federal organizations. Adoption remains uneven but is accelerating, and while some agencies continue to struggle with pilots, governance, and long-term implementation planning, the true challenge on the horizon is scaling these systems effectively and sustainably. Agencies must **optimize their AI implementations, scale responsibly, and prepare for the next wave of capabilities** that will reshape critical workflows in the coming years.

Organizations with mature data environments, strong oversight, and workforce readiness will be better positioned to adapt as AI capabilities evolve. Conversely, agencies that fail to operationalize AI effectively risk widening disadvantages against their faster-moving competitors and rival nations.

The Focus Has Shifted from Adoption to Optimization

Federal leaders now face a different challenge than they did even one year ago. **Adoption alone is no longer enough**. Agencies must simultaneously maximize impact while maintaining trust and efficiency. In addition, success increasingly depends not only on optimizing AI systems, but also on **equipping employees with the skills to use them effectively**.

Investment trends reinforce this shift. Gartner forecasts worldwide AI spending will reach $2.5 trillion in 2026, representing a 44% year-over-year increase.[1] Organizations are also increasing investments in AI infrastructure, data pipelines, and compute capacity to support larger-scale deployment.

The insights in this report are the product of ongoing collaboration between Everforth ECS experts working at the forefront of federal AI modernization, cross-referenced with learnings from industry and the commercial sector. Drawing on lessons from ***real-world implementations, observations from the field, and spirited discussions*** *about where the technology is headed next, the authors developed a shared perspective on the opportunities and risks agencies should be preparing for now.*

Building the Next Generation of Mission Systems

AI is transforming how technology itself is developed and deployed. Development teams are using AI-assisted coding, autonomous workflows, and AI-enabled testing pipelines to accelerate delivery timelines and reduce friction (84% of developers now use or plan to use AI coding tools).[2] These capabilities are helping agencies modernize more systems, and faster, than ever before.

AI is also reshaping cybersecurity. Defensive teams are increasingly relying on AI to identify vulnerabilities, accelerate patching, and counter adversaries who now have access to many of the same capabilities.

At the same time, agencies must balance speed with validation, explainability, security, and long-term maintainability. As AI-generated outputs become more common, the ability to verify accuracy and maintain trust becomes critical.

For federal leaders, the stakes continue to rise. AI implementation decisions made today will shape operational effectiveness and mission competitiveness for years to come.

The future of AI in government will be shaped by how effectively agencies operationalize trusted AI systems to achieve measurable mission impact.

1 Gartner 2 Uvik Software

Foundations

These are the non-negotiable capabilities and conditions necessary for data and AI success.

1. Evolving AI Ecosystems: From Tools to Orchestrated Systems

Federal agencies have largely moved beyond thinking of AI as a collection of standalone tools. Continued AI maturity depends on how effectively agencies integrate models, data, infrastructure, governance, and human workflows into **coordinated enterprise ecosystems**. That shift is already underway, but the pressure to operationalize AI in ways that are scalable, secure, and mission-aligned will continue to increase. (This direction aligns with the government's "AI Action Plan," which emphasizes adoption, interoperability, and accelerating the use of AI across government environments.)

Modern AI ecosystems now extend far beyond individual models. Agencies are building environments that combine agentic systems, semantic data layers (which act as translators to map technical jargon into familiar business terms), AI-assisted development, and automated workflows to improve speed and decision-making. As a result, **success increasingly depends on orchestration**. Systems must exchange context effectively, integrate with existing mission environments, and maintain accountability and oversight across the full AI lifecycle. Also, effective data management remains the foundation of these ecosystems because models, agents, and automated processes all depend on consistent access to trusted information.

Take the next steps:

- **Prioritize mission-aligned AI use cases** with measurable outcomes.
- **Build interoperable ecosystems** that connect data, models, infrastructure, and mission workflows.
- **Develop governance frameworks** that support speed accuracy, security, and long-term maintainability.
- **Build organization-wide AI competency** to improve adoption and oversight.

Agencies that treat AI as a disconnected collection of tools will struggle to scale implementation effectively. Organizations that invest in integrated, mission-ready ecosystems will be better positioned to adapt as AI capabilities continue to evolve.

> **"Collaboration, interoperability, and governance will define the next generation of federal AI success."**

Harkiran Kaur
Senior Product Manager, AI/ML

Foundations

2. Data Readiness Still Determines Success

AI systems are only as effective as the data environments that support them. As agencies expand AI adoption, **data readiness has become a defining factor in whether AI initiatives succeed, stall, or produce unreliable outcomes**.

The bottleneck is shifting. Organizations no longer struggle primarily with having too much data. Increasingly, the challenge is providing the right data, with the right context, to AI systems when and where it is needed.

And this challenge extends beyond data quality alone. Agencies must now manage data accessibility, governance, interoperability, and real-time availability across increasingly complex environments. Agentic systems and autonomous workflows depend on fast, trusted access to mission-relevant data. Delays, fragmented architectures, or inconsistent governance can **limit operational impact even when underlying AI models perform well**. Also, technical challenges are only part of the problem. Organizational barriers, data ownership concerns, and access restrictions often slow progress even when the underlying technology is available.

Federal leaders are already recognizing this shift. Recent Department of Homeland Security (DHS) CIO guidance has emphasized modern data architectures, such as data fabric and data mesh approaches, to improve enterprise-wide data accessibility, governance, and agility.[1]

Take the next steps:

- **Prioritize data hygiene** as an operational requirement.
- **Build architectures that support real-time access** across distributed mission environments.
- **Improve interoperability** between legacy systems, cloud platforms, and AI-enabled applications.
- **Develop semantic layers** that improve contextual understanding across data ecosystems.

> **Organizations that modernize their data environments now will be better positioned to scale AI securely, improve speed, and adapt as AI systems become more autonomous and context-aware.**

1 Homeland Security & Defense Forum (HSDF)

Foundations

3. AI Governance, Compliance, and the New Procurement Reality

Federal agencies are entering a new phase of AI adoption where governance, procurement, and oversight are becoming tightly interconnected. As AI systems move deeper into mission-critical environments, agencies must evaluate not only what AI can do but also **how those systems are acquired, governed, secured, and maintained over time.**

Federal guidance has reinforced this shift. In 2025, the Office of Management and Budget (OMB) released updated memoranda focused on accelerating federal AI adoption while strengthening governance, transparency, and procurement oversight.[1] The guidance places increased emphasis on **interoperability, vendor accountability, lifecycle monitoring, and reducing long-term vendor lock-in.**

Defense procurement priorities are also shifting rapidly toward **software speed, operational agility, and commercial-first acquisition strategies.** In 2025, the Secretary of War directed the Department of War (DoW) to prioritize modern software acquisition pathways, faster prototyping, and the expanded use of commercial contracting mechanisms, such as Commercial Solutions Openings (CSOs) and Other Transactions (OTs), to accelerate capability delivery to the warfighter.[2, 3]

Take the next steps:

- **Build governance into the full AI lifecycle**, including acquisition, deployment, monitoring, and retirement.
- **Prioritize interoperable platforms and architectures** that reduce long-term vendor dependency.
- **Strengthen explainability, auditability, and human oversight** for high-impact AI systems.
- **Align procurement strategies with readiness**, security requirements, and mission outcomes.

Governance must go beyond compliance. Agencies that integrate governance, procurement, and execution will be better positioned to navigate complex acquisition and oversight requirements.

1 Office of Management and Budget (OMB) 2 Department of War (DoW) 3 National Defense Magazine

Frontiers

These are the emerging shifts that will shape federal data and AI strategy over the next 12 months.

1. Mission-First AI: The Rise of Vertical Intelligence

The federal AI landscape is increasingly moving away from broad, one-size-fits-all AI platforms in favor of mission-specific systems designed around operational context, domain expertise, and measurable outcomes. This strategic pivot toward "Vertical Intelligence" (purpose built, tailored to specific domains, and designed around operational context) is **reshaping how agencies evaluate, deploy, and scale AI capabilities.** (Advances in AI-assisted development are making it easier to purpose-build solutions for specific mission needs. Agencies are becoming less dependent on heavily customized enterprise platforms and more capable of rapidly developing targeted capabilities that align with specific operational requirements.)

Deploying in high stakes environments like intelligence analysis, cybersecurity, or battlefield decision-making requires more than raw computational power. It demands deep environmental knowledge in which these AI tools operate. Generic models lack specialized data taxonomies and security protocols, hitting a context barrier to effective functioning when deployed. Agencies are realizing that forcing broad models to understand classified or tactical knowledge requires significant customization.

Take the next steps:

- Prioritize AI capabilities **aligned to specific mission outcomes and operational needs**.
- Develop architectures that **support domain-specific models, workflows, and governance requirements**.
- **Integrate AI systems directly into environments** rather than standalone analytic platforms.
- **Establish feedback loops** between operators, engineers, and mission owners to improve performance over time.

> "As AI and Agentic AI mature, federal agencies that combine cutting-edge technology with deep mission knowledge will thrive – that's **Vertical Intelligence.**"

Ketan Mane, Ph.D.
Director, Digital and AI Solutions

Frontiers

2. Agentic AI: From Assistants to Autonomous Systems

Agentic AI has already changed how organizations interact with data, software, and workflows. Many agencies now use AI assistants to support analysis, automate repetitive tasks, generate code, summarize information, and accelerate decision making. Current development trends are moving toward increasingly autonomous systems capable of **coordinating actions, maintaining context across environments, and operating with reduced human intervention**.

This evolution has driven rapidly growing interest in multi-agent architectures, contextual reasoning, and AI systems that can dynamically interact with enterprise data environments in real time. Knowledge graphs (maps of meaning, context, and semantic connections) and semantic data layers are becoming central because they help AI systems understand relationships and mission relevance rather than simply processing isolated data points. These technologies **improve decision quality, support explainability, and enable more adaptive workflows across complex environments**.

Federal agencies are already deploying these capabilities across cybersecurity operations, intelligence analysis, logistics coordination, and mission planning. At the same time, greater autonomy increases the importance of observability, oversight, and accountability. **Agencies must ensure autonomous systems remain transparent, secure, and aligned with mission intent**.

Take the next steps:

- **Develop architectures that support contextual reasoning** across distributed data environments.
- **Invest in semantic data layers and knowledge graph capabilities** to improve contextual understanding.
- **Establish governance frameworks** for autonomous and semi-autonomous AI systems.
- **Maintain human accountability** for high-impact mission decisions.

> **Leading AI systems already coordinate workflows, retrieve context dynamically, and execute tasks across environments. The next challenge is increasing autonomy while maintaining accountability, trust, and operational control.**

Frontiers

3. Speed to Impact vs. Trust and Validation

Federal agencies have already accelerated AI adoption across many fielded systems. Over the past year, agencies have expanded AI-assisted workflows, modernized procurement pathways, increased use of commercial AI capabilities, and shortened timelines for testing and deployment. The conversation is no longer centered on whether agencies can move faster. The focus now is **how to scale AI safely while maintaining deployment confidence**.

As deployment cycles accelerate, validation requirements are growing just as quickly. Faster deployment cycles require stronger monitoring, testing, and human accountability to maintain mission trust and reliability.

Take the next steps:

- **Build validation and governance** directly into AI development and deployment pipelines.
- **Prioritize measurable mission outcomes** over experimental activity.
- **Expand continuous monitoring** for AI-enabled systems and autonomous workflows.
- **Align operational speed** with security, explainability, and mission assurance requirements.

The agencies that gain the greatest advantage will not necessarily be the ones deploying AI the fastest. They will be the organizations that scale trusted AI systems effectively under real-world mission conditions.

Frontiers

4. Enabling the AI-Augmented Workforce

Federal workforces are already adapting to AI-enabled workflows across daily operations. Across government environments, personnel are using AI-assisted tools to support and speed up research, software development, documentation, and administrative activities. The next phase of workforce transformation will focus less on access to AI tools and more on how effectively agencies evaluate these tools, approve them, and **integrate them into daily operations, decision-making, and software development workflows**.

AI-assisted development is becoming one of the clearest examples of this shift. Coding assistants and AI-enabled development pipelines help teams accelerate modernization efforts, reduce repetitive work, and shorten delivery timelines.

AI Literacy Check

One measure of AI literacy is understanding which tools are best suited to which tasks.

- **Coding?** AI coding assistants.
- **Research?** Agentic research systems.
- **Image analysis?** Multimodal models.

Organizations that can consistently match the right tool to the right problem will be better positioned to capture value while avoiding unnecessary risk and inefficiency.

At the same time, workforce adoption remains uneven across many federal environments. Agencies continue to face challenges with effective adoption and workforce readiness. As AI-generated outputs become more common, employees must understand not only how to use these tools, but also how to **validate results, recognize limitations, and maintain human judgment in high-impact environments**.

The workforce itself is also evolving. Agencies are beginning to redesign processes around AI-assisted collaboration rather than isolated task automation. This will continue changing how teams collaborate, how software is developed, and how decisions are supported over the next several years.

Take the next steps:

- **Equip personnel** to work effectively with AI systems.
- **Establish validation and oversight processes for** AI-generated outputs and code.
- **Prioritize workforce training** focused on effective human-AI collaboration.
- **Establish validation and audit mechanisms** that support trust between humans, agents, and other AI systems.

> "Organizations that treat AI as a workforce multiplier rather than a standalone tool will be better positioned to **improve speed, reduce manual burden, and adapt** as AI capabilities continue to mature."

Mojgan Pedoeim
Director, Solutions Architecture

Fault Lines

These are the operational, structural, and strategic risks that can derail data and AI efforts.

1. Strategy Gaps and the Persistence of Pilot Purgatory

Federal agencies have made significant progress in AI adoption over the past several years. Many organizations now operate AI-enabled systems in production environments and support AI-assisted workflows across mission areas. Despite this progress, **many agencies still struggle to scale isolated successes into coordinated enterprise strategies**.

Pilot programs continue to create value, but too many remain disconnected from long-term planning. Agencies often prove that a capability works without fully addressing how it will be integrated into environments or sustained over time. Questions around ownership, security, sustainment, and workforce readiness are frequently addressed later in the process, creating friction when agencies attempt to scale successful pilots into broader deployments.

This challenge is becoming more visible as AI environments grow more complex. Agentic systems and autonomous workflows require coordination across leadership teams, acquisition offices, cybersecurity personnel, and operational stakeholders. Without a unified strategy, organizations can end up with **fragmented systems, duplicated investments, and competing priorities** across mission areas.

Pressure to deliver rapid results can also contribute to short-term decision making. Agencies are moving faster than they were even one year ago, but speed alone does not create lasting advantage. Long-term success depends on whether agencies can **connect experimentation to sustainable implementation strategies**.

Take the next steps:

- **Develop enterprise AI roadmaps** tied directly to mission priorities and measurable outcomes.
- **Define ownership models early** so systems can scale effectively over time.
- **Integrate governance, security, and sustainment planning** into deployment strategies from the beginning.
- **Focus on long-term integration,** not just initial demonstrations of capability.

> "Organizations that move beyond disconnected pilot activity **will be better positioned to scale trusted AI systems across enterprise and mission environments** as the technology continues to evolve."

Matthew McDonald
Vice President, Technology and Innovation

Fault Lines

2. Compute, Cost, and Infrastructure Constraints

AI adoption is increasing demand for compute storage, networking, and specialized infrastructure across federal environments. As agencies expand AI-enabled operations, infrastructure planning is becoming a major factor in **determining what systems can scale effectively and which initiatives stall under operational or financial pressure**.

Large-scale AI systems require significant investment in cloud infrastructure, GPUs, and large-scale model environments. In many cases, the cost of deploying and maintaining advanced AI systems still **exceeds the cost of the manual processes they are intended to support.** This is especially true for resource-intensive workloads involving large language models (LLMs), real-time inference, or autonomous workflows operating across distributed environments. Cost models are also evolving. As AI providers continue moving toward consumption-based pricing, agencies will need **greater visibility into usage patterns and long-term operating costs**.

Infrastructure complexity continues to grow, as well. Many agencies must support AI workloads across hybrid cloud environments, classified networks, edge systems, and legacy platforms that were not originally designed for AI-enabled operations. Performance, latency, interoperability, and cost management all become **more difficult as environments expand**.

As a result, agencies are increasingly prioritizing modular architectures, scalable data environments, and commercially proven deployment strategies that reduce unnecessary technical debt and infrastructure overhead. Experienced implementation partners can also help agencies accelerate deployment timelines, optimize resource utilization, and avoid costly overengineering during large-scale modernization efforts.

Take the next steps:

- **Prioritize infrastructure investments** that support long-term scalability.
- **Evaluate AI workloads** based on mission impact, operational efficiency, and total cost over time.
- **Modernize data pipelines and compute environments** incrementally to reduce disruption.
- **Improve observability across AI infrastructure** to monitor utilization, performance, and cost drivers.

> "As we move beyond the subsidized era of AI, organizations will need to **align their infrastructure strategies with scaling AI capabilities** to be successful and fiscally responsible."

Patrick Elder
Director, AI Solutions

Fault Lines

3. The AI Reality Gap: Navigating Slop vs. Value

As AI-generated content becomes more common, organizations are increasingly encountering what many in the technology community now refer to as "AI slop." The term generally describes low-quality or insufficiently validated AI-generated outputs that create **additional operational burden instead of meaningful value**.

This issue extends far beyond text generation. AI slop can appear in software development, documentation, analysis workflows, procurement artifacts, and reporting. AI-assisted coding tools, for example, can accelerate development significantly, but poorly reviewed AI-generated code may introduce **security vulnerabilities, inconsistent architecture patterns, technical debt, and maintainability challenges** that compound over time.

The long-term impact can be substantial. Organizations that prioritize speed without validation may eventually face rising sustainment costs, degraded system reliability, and large-scale remediation efforts to correct poorly governed AI-generated outputs embedded across enterprise environments.

Take the next steps:

- **Establish validation standards** for AI-generated code, analysis, and operational outputs.
- **Maintain human review processes** for high-impact systems and mission-critical workflows.
- **Monitor technical debt and long-term maintainability risks** associated with AI-assisted development.
- **Prioritize quality, traceability, and trust** alongside speed and productivity gains.

> "**AI is a force multiplier for an organization's expertise.** Organizations that lean into that by investing in both better AI and better expertise are the organizations best-positioned to grow long term value in the new AI-driven world."

Austin Amaya, Ph.D.
Director, AI and Analytics

Fault Lines

4. Trust, Security, and Explainability Risks

As AI systems become more integrated into mission operations, the consequences of failure become more significant. Federal agencies are increasingly relying on AI-generated analysis, autonomous workflows, and AI-assisted decision support in environments where errors or failures can directly affect mission outcomes.

Many agencies have already established governance frameworks and security controls for AI-enabled systems. However, the threat landscape continues to evolve alongside the technology itself. Risks such as **hallucinations, model manipulation, and unauthorized data exposure** are becoming more difficult to detect as systems grow more interconnected and autonomous.

Explainability is becoming more important in high-impact environments. Mission owners, operators, and oversight personnel must understand how AI systems arrive at conclusions. Systems that cannot provide sufficient transparency may struggle to gain long-term trust even when they deliver strong performance.

The increasing use of agentic systems and AI-assisted development pipelines further expands the challenge. As AI systems begin coordinating actions across enterprise environments, **organizations must maintain visibility into how decisions are made, what data is being accessed, and where operational responsibility ultimately resides.**

Take the next steps:

- **Strengthen continuous monitoring** for AI-enabled systems and autonomous workflows.
- **Implement explainability and auditability requirements** for high-impact environments.
- **Expand** red-teaming, adversarial testing, and validation processes for AI systems.
- **Establish validation, verification, and audit frameworks** for interactions between humans, agents, and autonomous systems.

> "**Trust must be central in AI deployment.** Agencies that balance speed with transparency and accountability will be better positioned to sustain AI adoption at enterprise scale."

Anthony Zech
Director, Data and AI

Fault Lines

5. The AI Talent and Literacy Gap

Federal agencies have made meaningful progress in expanding AI capabilities across technical and operational environments, but workforce readiness continues to **lag behind the speed of adoption**. Many organizations now have access to advanced AI systems without enough personnel who fully understand how those systems should be implemented, governed, validated, or sustained over time.

The challenge extends beyond hiring AI engineers and data scientists. Agencies increasingly need broader AI literacy across leadership teams, acquisition offices, cybersecurity personnel, software development teams, and mission operators. As AI becomes more integrated into daily workflows, **employees must understand how to evaluate AI-generated outputs, recognize limitations, and maintain human judgment in high-impact environments.**

Competition for experienced AI talent also remains intense. Commercial organizations continue offering faster hiring timelines, higher compensation, and more flexible career paths than many federal environments can currently support. As a result, **agencies are facing growing pressure to strengthen internal workforce development while continuing to modernize at speed.**

AI is also reshaping how work itself is performed. AI-assisted development tools, autonomous workflows, and decision-support systems are changing expectations across a variety of roles. Agencies that deploy advanced AI capabilities without investing in workforce readiness may **struggle to achieve long-term value**.

Take the next steps:

- **Expand AI literacy programs** beyond only technical teams.
- **Integrate AI expertise** directly into mission environments.
- **Train personnel to validate AI-generated outputs** and maintain effective oversight.
- **Develop career pathways** that improve retention and long-term workforce growth.

Long-term AI adoption success will depend on whether agencies can build workforces capable of managing and evolving AI systems responsibly as the technology continues to mature.

Conclusion

Across defense, homeland security, intelligence, and federal civilian environments, AI is already delivering value. The challenge now is sustaining that momentum and optimizing those implementations, while preparing for the **next generation of AI capabilities and innovation**.

Over the next 12 months, agencies will face growing pressure to scale AI responsibly while managing infrastructure demands, workforce readiness, and long-term trust. Organizations that establish strong data foundations, scalable architectures, and effective oversight models today will be better positioned to adapt as AI systems become more autonomous and embedded across enterprise workflows.

At Everforth ECS, we help federal agencies translate AI capability **into measurable mission impact**. Our teams support AI modernization efforts with expertise spanning data architecture, AI-assisted development, operational integration, cybersecurity, and enterprise-scale AI deployment.

If you would like to learn more about the topics discussed in this report or connect with our authors, we invite you to reach out and connect with us.

Matthew McDonald
Vice President, Technology and Innovation

Patrick Elder
Director, AI Solutions

Austin Amaya, Ph.D.
Director, AI and Analytics

Anthony Zech
Director, Data and AI

Ketan Mane, Ph.D.
Director, Digital and AI Solutions

Harkiran Kaur
Senior Product Manager, AI/ML

Mojgan Pedoeim
Director, Solutions Architecture

EverforthECS.com

From the boardroom to the battlefield, Everforth™ ECS® empowers our clients' missions, amplifies impact, and drives lasting results with frontier technology solutions in AI, cybersecurity, cloud, and data platforms. The company is focused on bringing the best of innovation to urgent mission needs across defense, intelligence, and federal civilian agencies. Everforth ECS maintains partnerships with leading providers of AI, cybersecurity, and more, bringing proven, mission-ready commercial technology to the government when and where it matters most.