ATARC | Qlik | gist360

TOP TAKEAWAYS

Building the Data Foundation for AI at Scale and the Future of Agentic AI in Government

A GIST 360 Federal Practitioner Roundtable White Paper

(Chatham House Rules)

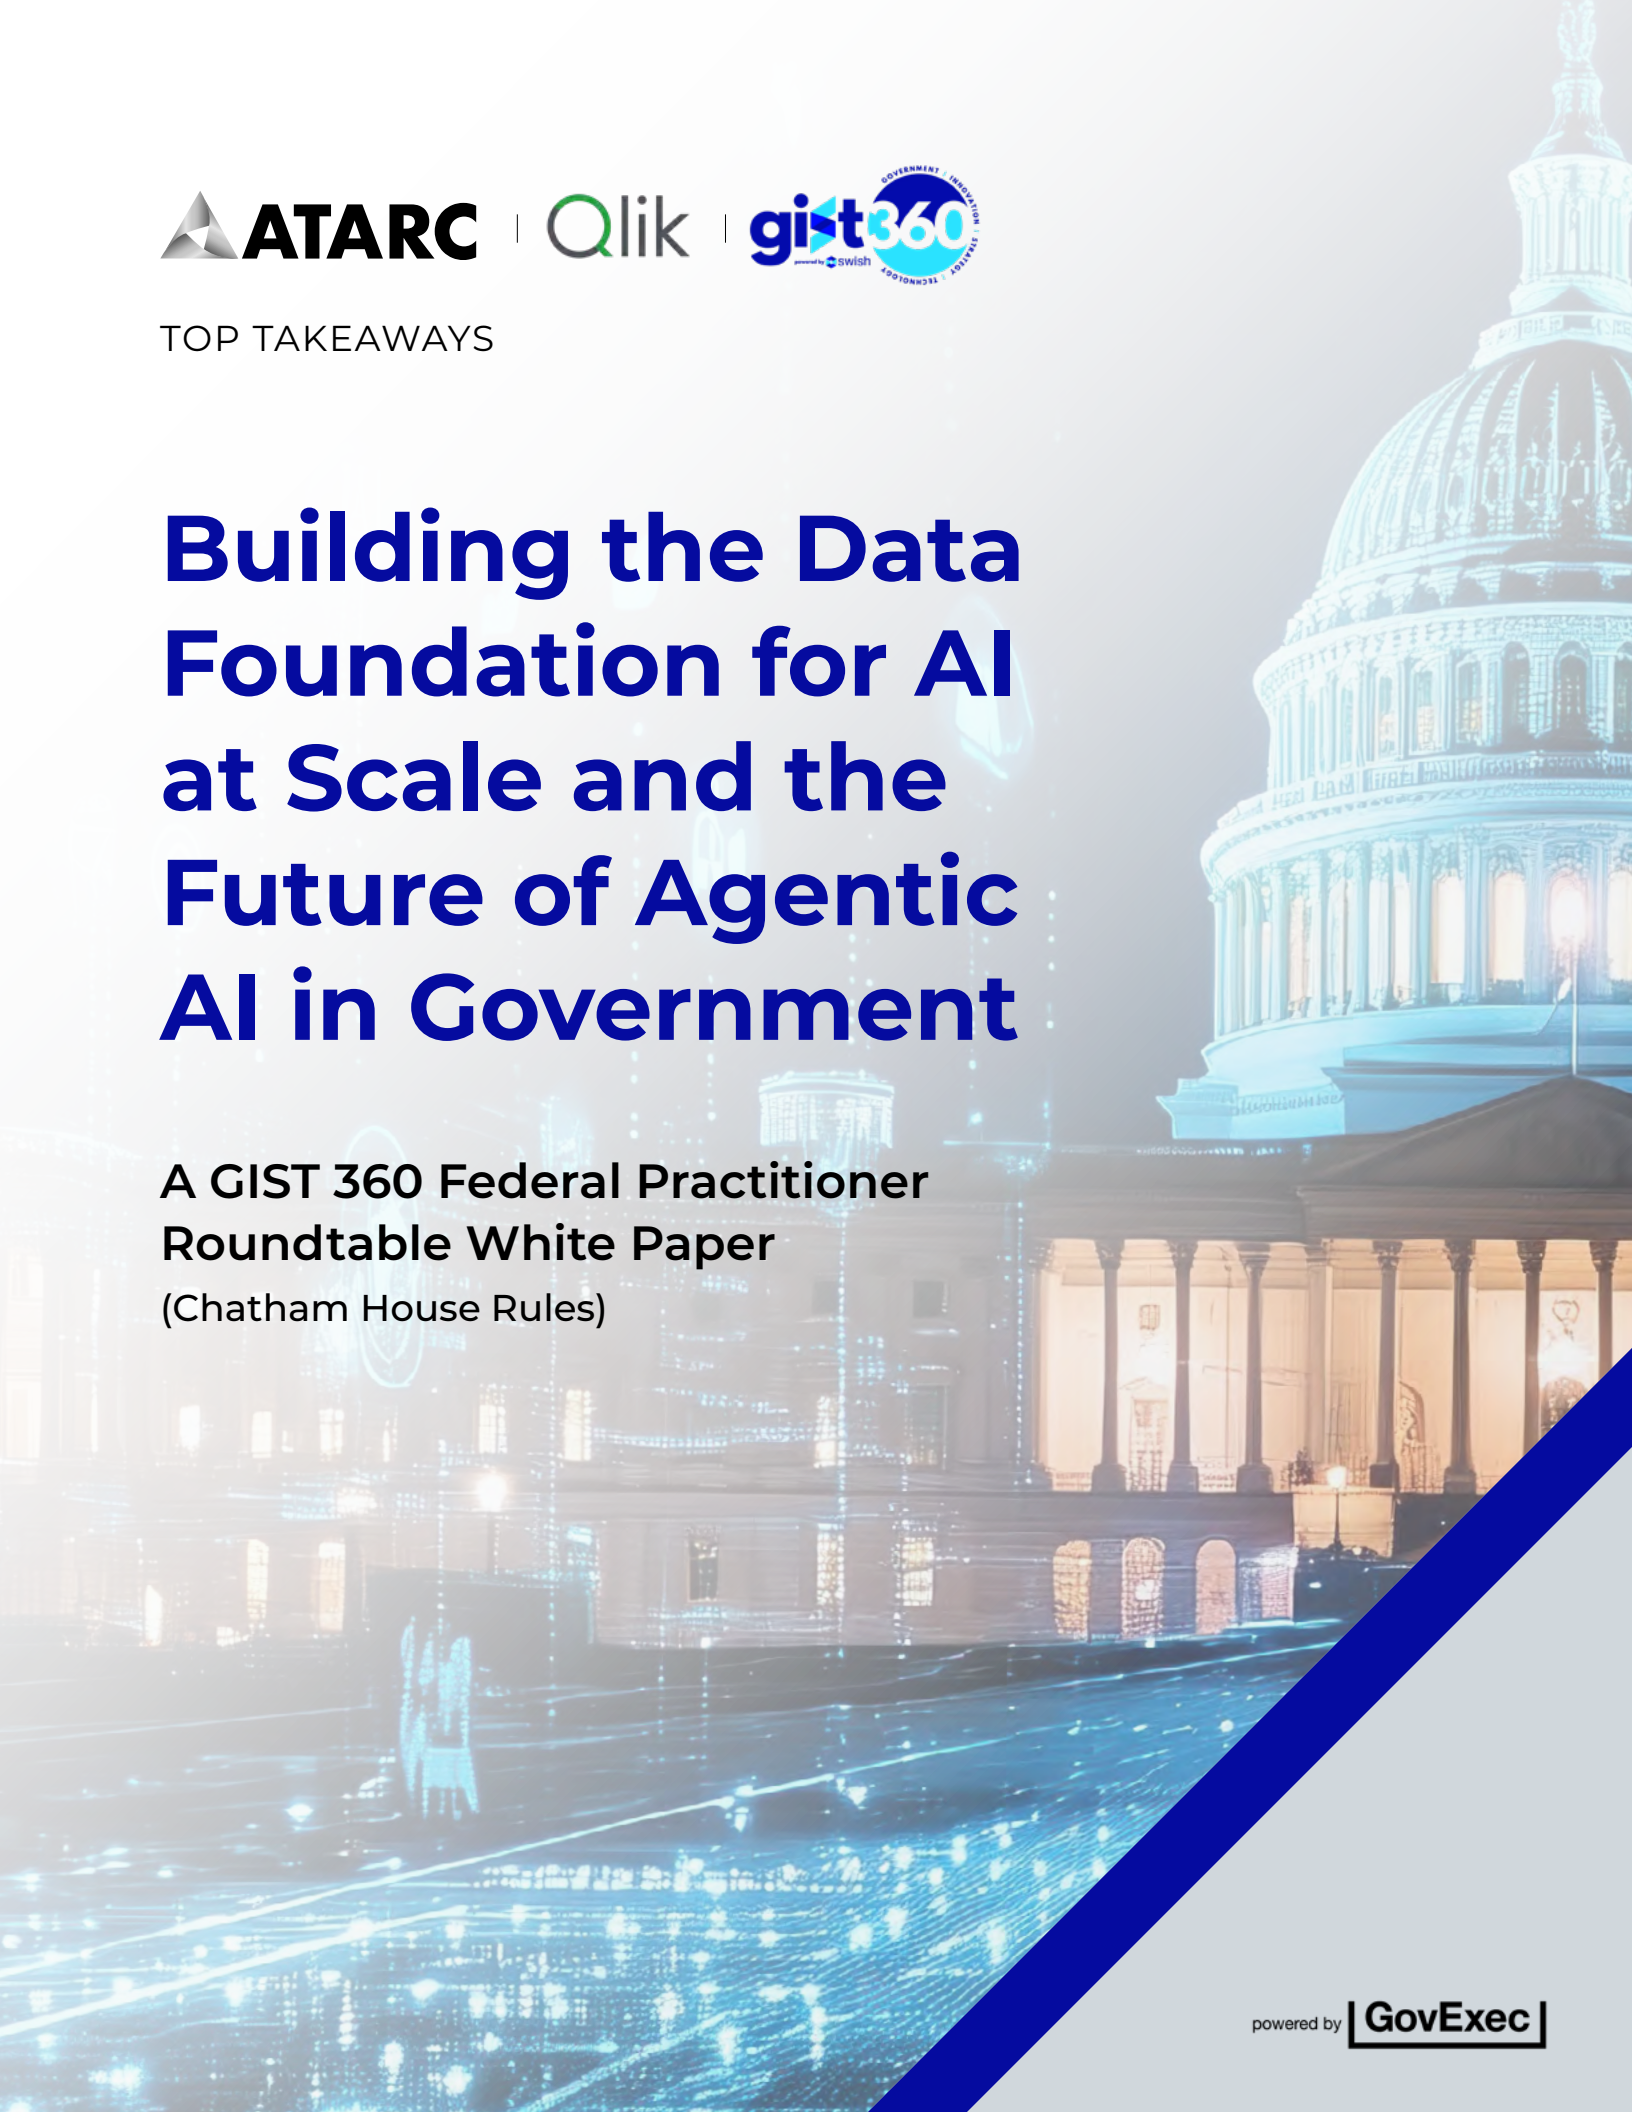

powered by GovExec

Executive Summary

Federal agencies are entering a decisive phase in their adoption of artificial intelligence. The early wave of experimentation, characterized by pilots, proofs of concept, and limited deployments, has largely validated the potential of AI. The challenge now is fundamentally different. Agencies must determine how to scale these capabilities across complex, mission-critical environments while maintaining trust, security, and operational integrity.

What emerged from this roundtable is a clear consensus: the primary barrier to scaling AI is not the technology itself. It is the condition, accessibility, and governance of data. Agencies are not struggling to access models or tools; they are struggling to ensure that the data feeding those systems is reliable, shareable, and fit for purpose.

At the same time, a new paradigm is beginning to take shape. Agentic AI, systems capable of acting autonomously, making decisions, and interacting across systems, introduces new demands that current architectures are not designed to support. This next phase will require not only better data, but a fundamentally different approach to how systems are designed, governed, and operated.

The discussion reflects a shift in focus from building AI to enabling it. The agencies that succeed will be those that treat data as infrastructure, not byproduct, and that design their systems for a future in which machines, not just humans, are active participants in mission execution.

From Pilots to Production: The Reality of Scale

In the early stages of AI adoption, agencies were able to demonstrate success within tightly controlled environments. These pilots often relied on curated datasets, clearly defined use cases, and limited user groups. Under those conditions, the performance of AI systems was often impressive.

However, as agencies attempt to scale these solutions, the underlying complexity of their environments becomes unavoidable. Data that appeared manageable in isolation is actually siloed and inconsistent when viewed across the enterprise. Systems that performed reliably in a pilot begin to struggle under the weight of increased volume, diversity, and interdependence.

This transition from pilot to production has proven to be one of the most difficult phases of the AI journey. It exposes gaps that were previously hidden, gaps in data integration, governance, and architecture. Agencies are discovering that scaling AI is not simply a matter of expanding existing solutions, but of rethinking the foundations on which those solutions are built.

In many cases, the lessons learned from pilots are less about what worked and more about what was missing. The controlled nature of pilot environments often masks the very challenges that define real-world implementation. As a result, agencies are increasingly recognizing the need to design for scale from the outset, rather than attempting to retrofit solutions after the fact.

What Data Readiness Really Means

The concept of data readiness has taken on new urgency in the context of AI. While it is often discussed in terms of availability, practitioners emphasized that availability alone is insufficient. Data must also be trustworthy, accessible, and meaningful within the context of its use.

Trust is a central concern. Agencies must be confident not only that their data is accurate, but that they understand where it came from, how it has been transformed, and whether it can be relied upon in decision-making. This requires a level of visibility and transparency that many organizations do not yet possess.

Accessibility presents another challenge. Data may exist within an agency, but that does not mean it is available to those who need it. Security boundaries, system incompatibilities, and organizational silos all contribute to an environment in which data is technically present but practically unusable.

Timeliness is equally important. In many mission scenarios, data that is even slightly outdated can lead to suboptimal or incorrect decisions. As AI systems increasingly operate in near real time, the expectation for current or streaming data becomes more pronounced.

Finally, context is essential. Data without meaning—without metadata, definitions, and lineage—cannot be effectively used by either humans or machines. As agencies move toward more advanced AI capabilities, the importance of contextualized data becomes even more critical.

Breaking Down Silos: The Cultural Dimension of Data

One of the most persistent barriers to data readiness is not technical, but organizational. Federal agencies have long operated within structures that encourage localized ownership of data. Programs, offices, and mission areas often maintain control over their own datasets, with limited incentives to share them.

This model is increasingly incompatible with the needs of AI. Effective AI systems require access to diverse and integrated data sources. When data is confined within silos, the potential value of AI is significantly diminished.

Shifting away from this model requires more than new technology. It requires a change in mindset. Data must be viewed not as a resource owned by individual components, but as a shared asset that supports the broader mission.

This transition is not without challenges. Concerns about security, privacy, and control are valid and must be addressed. However, the discussion made clear that these challenges can be mitigated through governance, transparency, and trust-building.

In practice, progress often begins with small, targeted efforts. By identifying use cases that benefit multiple stakeholders, agencies can demonstrate the value of shared data and build momentum for broader adoption. Over time, these successes can help shift organizational culture and establish new norms around data sharing, for example through enterprise-wide data catalogs, marts, warehouses or lake houses.

The Imperfection of Data and the Need for Pragmatism

A recurring theme throughout the discussion was the recognition that federal data is inherently imperfect. Errors in data entry, inconsistencies across systems, and gaps in completeness are common and, in many cases, unavoidable.

Rather than attempting to eliminate these issues entirely, agencies are adopting a more pragmatic approach. The focus is shifting from achieving perfect data to ensuring that data is fit for purpose. This involves understanding the limitations of the data, assessing the risks associated with its use, and making informed decisions based on that context.

This approach requires a high degree of transparency. Users of data, whether human or machine, must be aware of its strengths and weaknesses. Metadata, lineage, quality indicators, and trust score play a critical role in providing this visibility.

It also requires a shift in expectations. AI systems must be designed to operate within environments where data is not perfect. In more mature environments, they have adopted Medallion Architectures where a color schema represents different layers of quality. For example bronze is raw data exactly as it arrives, silver is cleaned and validated data ready for self-service analytics, and gold represents enriched data with additional information and business logic

Ultimately, the goal is not to eliminate imperfection, but to manage it effectively.

Governance, Lineage, and the Question of Trust

As AI systems become more integrated into mission operations, the question of trust becomes increasingly important. Decisions made by AI, whether recommendations or autonomous actions, must be explainable and defensible.

This places new demands on data governance. Agencies must be able to trace the origins of the data used by AI systems, understand how it has been processed, and verify its integrity. Without this capability, trust in AI outputs will remain limited.

Lineage and provenance are central to this effort. Knowing where data comes from and how it has been transformed provides the foundation for accountability. It also enables agencies to identify and address issues when they arise.

Auditability is another critical component. As AI systems begin to operate with greater autonomy, it is essential to maintain records of how decisions are made. This not only supports oversight, but also enables continuous improvement by providing insights into system behavior.

These capabilities are not optional. They are essential for ensuring that AI systems can be used safely and effectively in high-stakes environments.

Security and Integrity in the Age of AI

The integration of AI into federal systems introduces new security considerations. While traditional cybersecurity concerns remain relevant, AI creates additional vulnerabilities that must be addressed.

One of the most significant risks is data poisoning. If the data used to train or operate AI systems is compromised, the outputs of those systems can be misleading or incorrect. This risk is particularly concerning in environments where decisions have significant consequences.

Maintaining data integrity is therefore a top priority. This includes validating data sources, monitoring for anomalies, and ensuring that data has not been altered in unauthorized ways. It also involves securing the pipelines through which data flows, from ingestion to processing to output.

In addition, agencies must be vigilant in monitoring the behavior of AI systems themselves. Models can drift over time as new data is introduced, leading to changes in performance. Detecting and addressing this drift is essential for maintaining reliability.

These challenges underscore the need for a comprehensive approach to security—one that encompasses not only systems and networks, but also the data and models that drive AI.

The Emergence of Agentic AI

While much of the discussion focused on current challenges, there was also significant attention given to the future of AI in government. In particular, the concept of agentic AI represents a major shift in how these technologies are used.

Unlike traditional AI systems, which provide outputs for human interpretation, agentic systems are designed to take action. They operate continuously, make decisions in real time, and interact with multiple systems and data sources.

This shift has profound implications. It changes the role of humans from direct operators to overseers and exception handlers. It also increases the importance of data quality, governance, and system design.

Agentic AI systems do not interpret data in the same way humans do. They rely on structured inputs, clear definitions, and consistent formats. Ambiguity, which humans can often navigate, becomes a significant challenge for machines.

As a result, the requirements for data and systems become more stringent. Agencies must ensure that their environments are not only capable of supporting AI, but are optimized for it.

Designing for an AI-Driven Future

One of the most important insights from the discussion is that most existing systems were designed for human users. They assume a workflow in which individuals review information, apply judgment, and make decisions.

Agentic AI challenges this model. In a machine-driven environment, decisions are made at the point of data consumption, often without human intervention. This requires systems that are fundamentally different in how they present and manage information.

Applications and data must be structured in ways that allow for easy and efficient interactions with AI agents and bots. Definitions must be explicit. Relationships between data elements must be clearly defined. Access must be available in real time, often through APIs and other programmatic interfaces such as the Model Context Protocol (MCP) or Agentic AI Browsers.

Equally important is the ability to observe and understand system behavior. As machines take on a more active role in decision-making, it becomes essential to capture how those decisions are made. This includes not only the data used, but also continuous evaluation of the output, performance, and costs so they can be optimized over time..

Designing for this future will require agencies to rethink their architectures, moving away from human-centric models toward systems that support both human and machine interaction.

Conclusion

The path to AI at scale in government runs through data. While advances in models and algorithms continue at a rapid pace, their impact will be limited without a strong data foundation.

This foundation must address not only technical challenges, but organizational and cultural ones as well. It requires a commitment to data sharing, governance, and continuous improvement. It also requires a willingness to adapt to new paradigms, including the rise of agentic AI.

The agencies that succeed will be those that recognize this shift and act accordingly. By investing in data as a strategic asset and designing systems for a machine-driven future, they will be positioned to unlock the full potential of AI.

The stakes are high, but so is the opportunity. AI has the potential to transform how government operates, enabling more efficient, effective, and responsive services. Realizing that potential will depend not on the sophistication of the technology, but on the strength of the foundation beneath it.

About This Report

This white paper reflects insights from a federal practitioner roundtable conducted under Chatham House Rules. The perspectives presented capture shared themes and experiences across agencies, without attribution to specific individuals or organizations.