

NOVEMBER 2025

# Enabling Enterprise-grade AI with the NetApp Data Platform

Simon Robinson, Principal Analyst

**Abstract:** The promise of AI to drive business transformation is clear, but mainstream enterprises are struggling to achieve production-ready deployments that can deliver real ROI. As organizations increasingly identify a range of challenges in preparing AI-ready data, attention is turning to modern data infrastructure solutions that meet the high expectations of enterprise customers. In response, NetApp has unveiled its AFX AI portfolio as a new addition to its comprehensive, enterprise-grade data platform that spans an organization's entire data estate, accelerating data pipelines end to end with powerful performance, proven resiliency, and industry-leading integrated data security.

## Data Infrastructure Challenges Threaten AI's Promise in the Enterprise

Of all the technology innovations that have transformed our lives over recent decades, the potential of AI to drive foundational, meaningful improvements in the way we live and work is perhaps the greatest and most exciting. Yet, this excitement is already being tempered with a hard reality: Applying the undoubted vast potential of emerging AI technologies to individual businesses to drive real results, and ROI, is difficult.

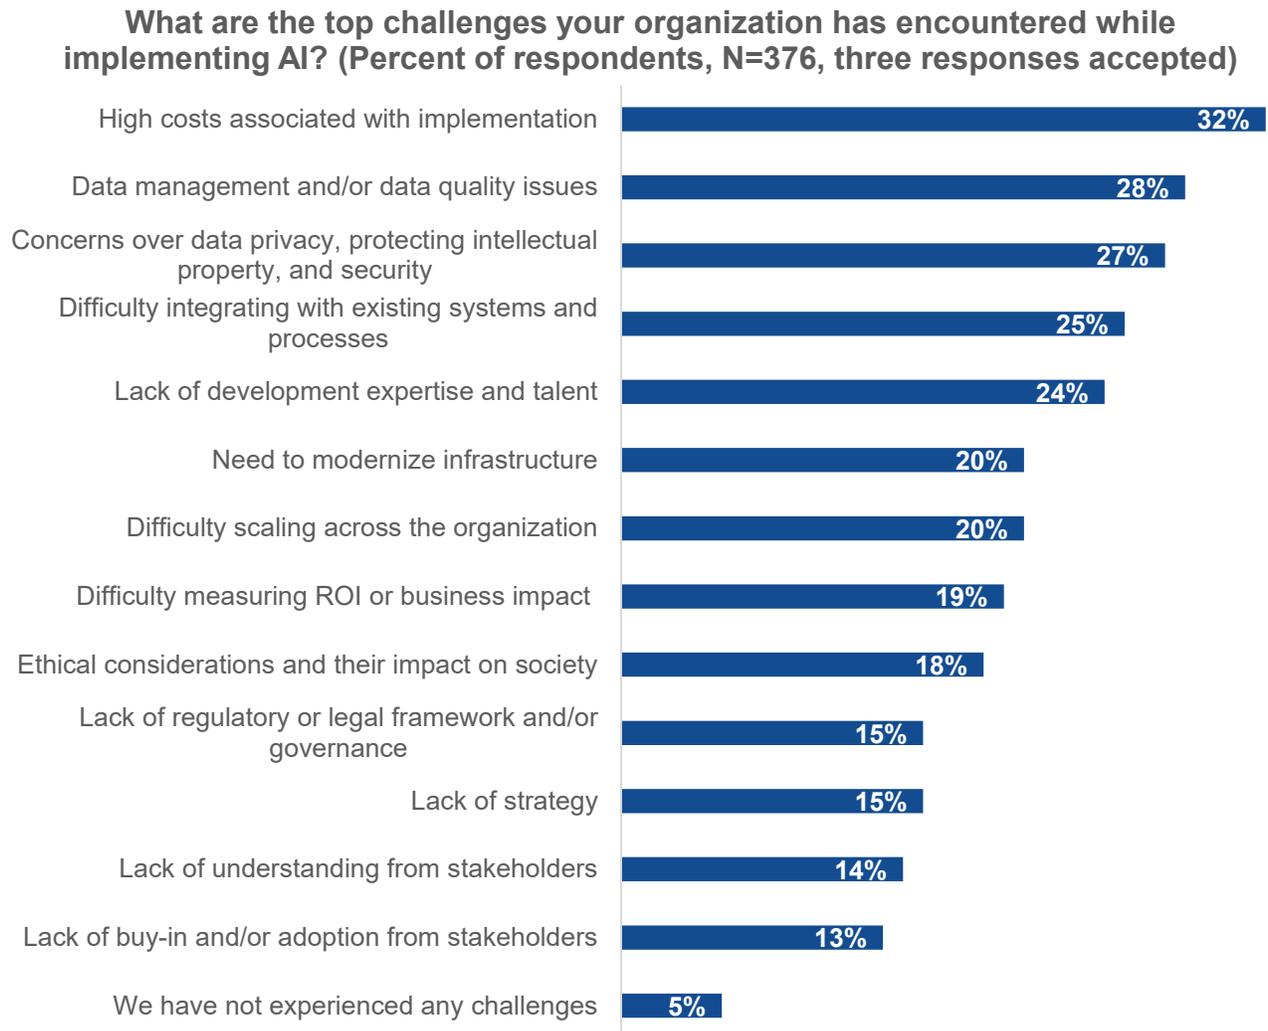
Even at this early stage, it is clear many organizations are facing substantial challenges that threaten their ability to fully take advantage of these new technologies. For example, research from Enterprise Strategy Group found almost all organizations (a massive 97%) have experienced challenges moving their AI efforts from development (proof of concept, POC) to production, with 62% reporting they have faced moderate to extensive challenges here.<sup>1</sup>

What's more, many of the key challenges are data related. For example, for AI implementations overall, data quality/management challenges are the second most commonly cited challenge, behind only cost. Additionally, concerns over data privacy ranked third (see Figure 1). For AI infrastructure specifically, the top two challenges cited were security and compliance, followed by data management and quality.<sup>2</sup>

Therefore, it's clear that a successful enterprise AI strategy requires a strong data foundation that can address a range of data management and data privacy/security issues.

<sup>1</sup> Source: Enterprise Strategy Group Research Report, [Navigating Build-versus-buy Dynamics for Enterprise-ready AI](#), January 2025.

<sup>2</sup> Ibid.

**Figure 1. Top AI Implementation Challenges**

Source: Enterprise Strategy Group, now part of Omdia

## Reality Check: Enterprise AI Is Different Than Hyperscale AI

A huge part of the challenge of implementing AI is that the overall AI narrative has been dominated by a small number of organizations developing foundational large language models (LLMs) and deploying them at hyperscale. Though these innovations have been fundamental in democratizing and unleashing the power—and potential—of AI, the reality is that deploying advanced AI technologies within the enterprise setting is a very different proposition.

On one hand, hyperscale AI is characterized by access to massive, open datasets, unconstrained experimentation, and seemingly endless budgets; it is also primarily focused on building, training, and fine-tuning the large models that are increasingly in common usage. By contrast, the typical enterprise looking to begin its AI journey is not interested in training LLMs; though most will start their AI projects using established LLMs, the value for them is going to be around retrieval-augmented generation (RAG), inferencing, and agentic use cases. In other words, organizations' main concern is how to effectively and efficiently build data pipelines using curated datasets from within their own teams. Critically, enterprises also have a range of additional considerations, including complexity,

legacy IT, compliance/governance, security, and, of course, cost, that might not be as relevant to hyperscale AI providers.

Despite this, the current AI solution landscape sometimes feels as if these “real-world” considerations haven’t been taken into account when, of course, they are critically important. For enterprises, data security and privacy is non-negotiable, while performance, simplicity, and data protection are similarly essential. Enterprises also need AI infrastructure that fits within existing skill sets, integrates with current capabilities, and enforces strict governance.

Let’s take a closer look at some of the specific challenges mainstream enterprises are facing as they ramp up their AI initiatives.

## Data Fragmentation Inhibits Foundational Data Access

It’s well understood that any AI strategy needs to access the right data to be effective. However, the modern enterprise is highly distributed, and so is its data—on-prem, in the public cloud, and at the edge. This data is also often trapped inside an assortment of poorly integrated storage silos at each location; more than three-quarters of IT leaders said that managing data across a hybrid cloud environment is a substantial challenge.<sup>3</sup> This fragmentation makes it very difficult to gain a global view of data, resulting in multiple potential issues. For example, if data/AI scientists can’t easily access or understand all an organization’s data, this will limit their ability to leverage it in their AI workflows (e.g., to deploy RAG or fine tune a model).

Data fragmentation introduces risks, including that wrong, inappropriate, or sensitive data will inadvertently be used to train or fine-tune a model. Moreover, any data used in an AI workflow needs to be protected from ever-evolving cybersecurity threats such as ransomware. According to recent Enterprise Strategy Group research, 83% agreed that success in AI is impossible if data isn’t secured and protected.<sup>4</sup> Both of these issues drive up cost and complexity, inhibiting an organization’s ability to effectively move from POC to production or negatively impacting ROI.

Some industry solutions advocate for the creation of dedicated storage environments for AI data. While this may simplify certain aspects of the AI data process, adding another silo to the environment may actually add to the data fragmentation issue. Conversely, AI solutions that integrate with the existing data and storage infrastructure may appeal to organizations already struggling with data fragmentation and related complexity.

## AI Tooling Complexity Slows the Data Pipeline

Compounding AI challenges is additional complexity at the data layer: 60% of organizations have experienced greater complexity within their data environment because of AI workload integration.<sup>5</sup> Moreover, 71% cited major challenges around effectively integrating storage with data pipelines for AI.<sup>6</sup> One of the key challenges here is a cluttered collection of AI tools. The AI tool set landscape has evolved rapidly, but this has led to organizations adopting a dozen or more point tools, often in a haphazard mishmash that inhibits their ability to quickly bring curated datasets from across their entire data estate to data scientists.

A lack of integration or awareness between AI tools and the underlying storage environment threatens to further compound this complexity, as well as drive “data bloat,” as the same tasks could be performed multiple times across the same dataset.

---

<sup>3</sup> Source: Enterprise Strategy Group Research Report, [The Critical Role of Storage in Building an Enterprise AI Infrastructure](#), September 2025.

<sup>4</sup> Source: Enterprise Strategy Group Research Report, [IT Transformed: Inside the Convergence of Hybrid Cloud and AI](#), July 2025.

<sup>5</sup> Source: Enterprise Strategy Group Complete Survey Results, [Data Readiness for Impactful Generative AI](#), April 2025.

<sup>6</sup> Source: Enterprise Strategy Group Research Report, [The Critical Role of Storage in Building an Enterprise AI Infrastructure](#), September 2025.

## Limitations of Legacy and Alternative Data Storage Approaches

AI workloads have high demands on the underlying infrastructure, and this includes storage. According to Enterprise Strategy Group research, 70% of organizations said storage challenges are a significant barrier to AI success.<sup>7</sup> Though the challenges vary, performance is a particular issue. For example, GPUs are expensive and have very high performance demands, so having GPUs sit idle while waiting for storage to feed them data is highly inefficient.

Meanwhile, AI-centric data transformations (such as vectorization) can lead to ballooning data volumes; some organizations have experienced significant data bloat as they implement AI. With 87% reporting that AI will or is already driving substantial data growth at their organizations,<sup>8</sup> this seems to be a common experience that could substantially drive-up costs.

AI use cases are also rapidly evolving. It can be difficult to anticipate what infrastructure requirements will be required over the next two years' time, so it's crucial organizations have an agile and flexible data infrastructure that can rapidly scale as their requirements grow and change. Critically, this should include the ability to scale storage compute independently from capacity.

At the same time, storage for AI cannot be a science project that requires high levels of skill and specialization. Solutions need to take advantage of existing skill sets while also taking into consideration the existing data estate and leveraging existing processes and well-understood enterprise practices around resiliency, security, and data management. Once again, while alternate approaches may advocate for the deployment of an additional silo for dedicated AI storage, this may increase the overall management burden, requiring additional specialist skills to manage dedicated silos that may not be integrated with the wider storage and data environment.

Taken together, these issues should prompt IT leaders to consider their own AI journey and challenge themselves to better understand the fundamentals of their own efforts to move from POC to production within their own enterprise setting. As part of this, they should ask whether their current data infrastructure is fit for purpose in the AI era and whether it enables or inhibits a secure, compliant, and performant AI data strategy. Additionally, when assessing their options, IT leaders should also consider whether they want to build a brand-new AI environment from scratch or whether an evolutionary approach that enables them to embrace AI-driven workloads within the context of their existing skill sets, processes, and data environment is more appropriate.

## Introducing the NetApp AFX AI portfolio – An Enterprise-grade Data Foundation for the AI Era

As a leading provider of intelligent data infrastructure, NetApp already offers a comprehensive data platform designed to help customers optimize for AI workloads. However, those efforts are now taking a significant step forward through a range of new innovations that collectively form the NetApp AFX AI portfolio (see Figure 2).

Rather than taking a piecemeal approach to accelerating AI workloads, NetApp has adopted a comprehensive approach designed to support an organization's AI journey in a much more holistic fashion. Designed to simplify the life of infrastructure and data professionals alike, as well as deliver bullet-proof security, reliability, and operational efficiency, the NetApp data platform is billed as the enterprise-grade data platform for AI.

---

<sup>7</sup> Ibid.

<sup>8</sup> Ibid.

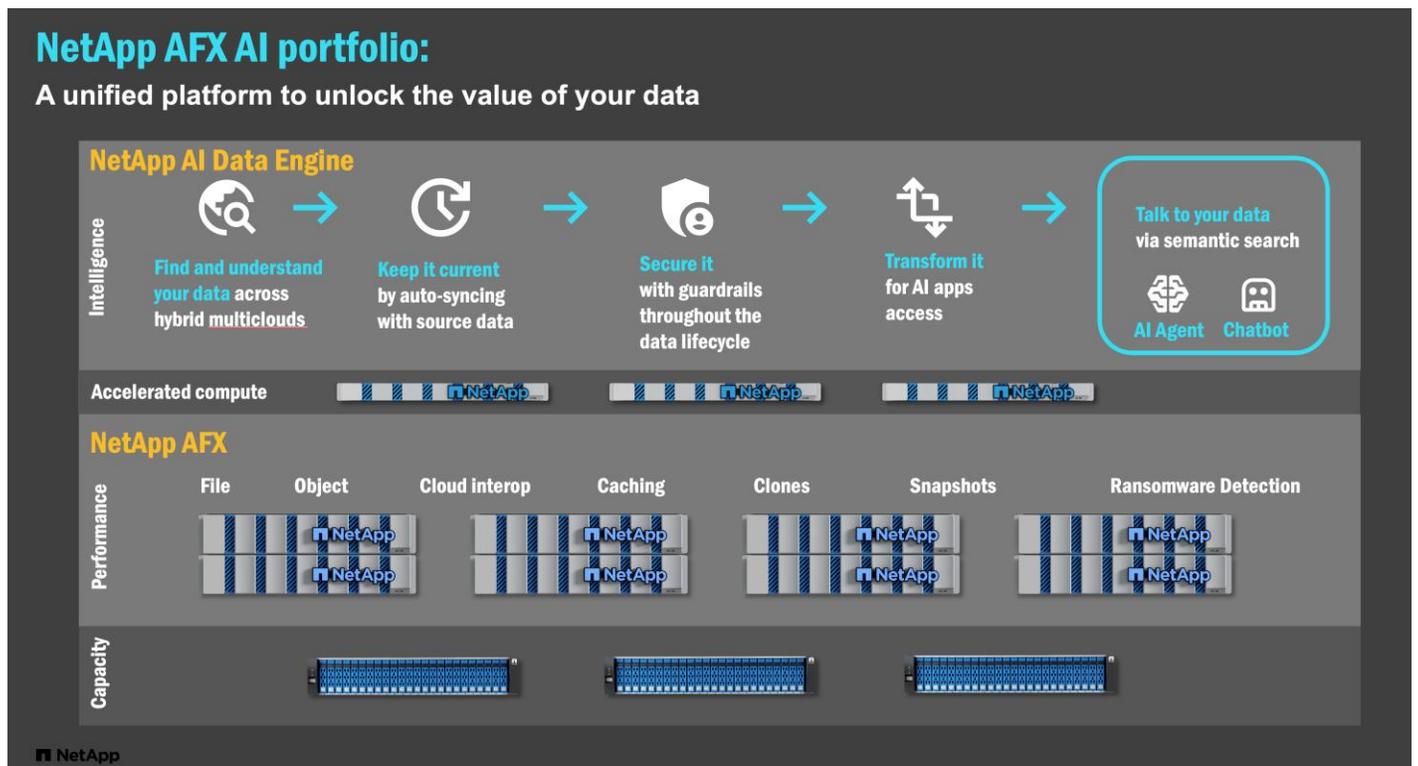
The data platform delivers the following capabilities:

- **Enterprise resilience and performance.** Built on NetApp's decades-long experience of storing a significant proportion of the world's unstructured data, NetApp AFX offers modern, disaggregated storage with proven enterprise-grade high availability, parallel file system performance without the headaches, and bulletproof security.
- **Unified data foundation.** Leveraging NetApp's unique ability to unify a customer's data environment—be that on premises or across any of the major hyperscale cloud (Google Cloud, AWS, Azure)—the NetApp data platform delivers a single view of the data estate, making it instantly available wherever and whenever needed. A global metadata catalog and unified namespace provide visibility, and tools such as integrated zero-copy caching and instant data clones end the copy sprawl nightmare, while secure multi-tenancy enforces logical separation and data granularity across common infrastructure.
- **Accelerated data pipeline.** Rather than a cluttered drawer of AI tools, the NetApp AI Data Engine delivers a streamlined, real-time method to instantly bring curated datasets to data scientists to rapidly drive insights and agentic action, end to end. An embedded vector database increases data processing efficiency while reducing data bloat, policy-based guardrails automatically identify and protect sensitive data to eliminate compliance blind spots and boost confidence in regulatory alignment, and support for the NVIDIA AI Enterprise Platform enables zero-copy integration with major cloud AI platforms and popular open source AI tools.

As noted, the NetApp AFX AI portfolio includes a range of new capabilities that build on NetApp's multi-decade experience of delivering a range of performant, reliable, and secure unstructured data solutions that enterprises globally run their businesses on. Key aspects of the portfolio include:

- **NetApp AFX.** A modern disaggregated storage system running the NetApp ONTAP operating system, specifically optimized for AI workloads. Built on a disaggregated architecture that enables independent scaling of storage compute and capacity, AFX delivers scalable, optimized storage performance without overprovisioning, and it's backed by the enterprise-grade resilience and security features NetApp is known for.
- **NetApp AI Data Engine.** A set of powerful features that accelerate the enterprise AI data pipeline, including:
  - **Metadata Engine.** Automatically creates a global, structured, interactive view of data, enabling customers to easily explore and understand their data for optimal management. Lightweight APIs enable optimized metadata queries.
  - **Data Sync:** Eliminates redundant data and streamlines hybrid cloud AI, ensuring data is always current.
  - **Data Guardrails:** End-to-end protection and privacy for AI workloads. Continuously classifies and anonymizes and governs sensitive data, automatically. This ensures AI models only utilize compliant data (i.e., sensitive data is never vectorized), reducing risk, aligning with enterprise security policies, and boosting confidence for data scientists and engineers. In turn, this reduces breach risks and penalties in accordance with customer policies, enabling highly secure and compliant AI-powered apps.
  - **Data Curator:** Transforms and optimizes data for generative AI (GenAI) apps such as RAG. NetApp's vectorization engine powers low-latency retrieval for GenAI apps such as RAG, using compressed embeddings (reducing data bloat by up to 10x, according to NetApp) and re-ranking with LLMs. This enables scalable, cost-efficient GenAI across hybrid environments.

Figure 2. NetApp AFX AI portfolio



Source: NetApp

## The NetApp AFX AI portfolio Drives Transformational Value Across a Variety of Data Roles and Use Cases

With the NetApp AFX AI portfolio, customers can boost not just their infrastructure performance but also their end-to-end pipeline efficiency. By eliminating redundant data movement, automating sync, and integrating with AI tools, NetApp reduces latency and speeds time to value. As a result, data scientists and data engineers can begin to deploy production-ready AI workloads without limitation and without storage or data access acting as a barrier. Example benefits include:

- Enabling data scientists to use their AI tool of choice for RAG workloads or fine-tune a model from a data catalog that spans their entire data estate—without them needing to know where the data is stored.
- Enabling data scientists to create curated datasets of high-quality, compliant data, then process the data and integrate it with their cloud AI tools—efficiently and all in one place.
- Capturing stateful training data with efficient point-in-time copies that enable data scientists to go back and analyze data in its original state, letting them understand in detail a model's decisions. Crucially, this is enabled regardless of whether the training data resides on premises or in any major public cloud.
- All built on a fast and efficient storage foundation that can independently scale performance and capacity based on evolving requirements.

## Conclusion

Enterprises of all sizes are looking to embrace the potential of AI technologies within their environments, but to be effective, they need to be confident that these solutions meet their needs for enterprise-grade resilience, security, and efficiency. Success here isn't just about building a better mousetrap; it's about developing innovative solutions that embrace new opportunities in a manner that makes sense for enterprises in the real world.

In this sense, the timing of the NetApp AFX AI portfolio couldn't be better; NetApp has built on its heritage as a proven intelligent data infrastructure solution and is taking it into the AI era with the NetApp data platform. As a result, customers stand to not just boost their infrastructure performance for AI but also accelerate their AI data pipelines end to end, with confidence. Any organization contemplating its AI journey from POC to production is advised to take a closer look at NetApp AFX AI portfolio.

To learn more about NetApp AFX AI portfolio, please visit <https://www.netapp.com/afx/>.

©2025 TechTarget, Inc. All rights reserved. The Informa TechTarget name and logo are subject to license. All other logos are trademarks of their respective owners. Informa TechTarget reserves the right to make changes in specifications and other information contained in this document without prior notice.

Information contained in this publication has been obtained by sources Informa TechTarget considers to be reliable but is not warranted by Informa TechTarget. This publication may contain opinions of Informa TechTarget, which are subject to change. This publication may include forecasts, projections, and other predictive statements that represent Informa TechTarget's assumptions and expectations in light of currently available information. These forecasts are based on industry trends and involve variables and uncertainties. Consequently, Informa TechTarget makes no warranty as to the accuracy of specific forecasts, projections or predictive statements contained herein.

Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of Informa TechTarget, is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact Client Relations at [cr@esg-global.com](mailto:cr@esg-global.com).

---

### About Enterprise Strategy Group

Enterprise Strategy Group, now part of Omdia, provides focused and actionable market intelligence, demand-side research, analyst advisory services, GTM strategy guidance, solution validations, and custom content supporting enterprise technology buying and selling.

 [contact@esg-global.com](mailto:contact@esg-global.com)

 [www.esg-global.com](http://www.esg-global.com)