

NetApp Focuses on Enterprise-Grade AI Infrastructure with New AFX Storage System and NVIDIA-Fueled AI Data Engine

October 30, 2025

By: [Carol Sliwa](#), [Natalya Yezhkova](#), [Jasdeep Singh](#), [Phil Goodwin](#)

IDC'S QUICK TAKE

NetApp strengthened its intelligent data infrastructure strategy for AI workloads with the launch of new products and capabilities at its October [INSIGHT 2025](#) conference in Las Vegas. The company introduced an all-flash AFX storage system designed to independently scale performance and capacity and an AI Data Engine (AIDE) that integrates its ONTAP operating system and data services with NVIDIA AI Enterprise software and runs on GPU-based accelerated compute nodes in the same hardware rack. Additional announcements included the expansion and enhancement of NetApp's public cloud, cyber-resilience, and security capabilities.

EVENT HIGHLIGHTS

The most significant announcements at NetApp INSIGHT 2025 focused on enterprise storage infrastructure and data services for AI workloads. The new all-flash AFX storage system introduces a disaggregated storage architecture option designed to separately scale compute and storage resources to meet the performance and capacity demands of AI workflows while providing the enterprise capabilities of ONTAP. NetApp says an AFX cluster will be able to scale to a maximum of 128 storage controllers and 52 NX224 NVMe storage enclosures, for throughput up to 4TBps and effective capacity of more than an exabyte. ONTAP-based AFX storage is also certified for NVIDIA DGX SuperPOD use. Customers have the option to add up to 10 AFX DX50 data compute nodes, equipped with AMD Genoa CPU and NVIDIA L4 GPU, to power the AIDE software and a new metadata engine to index data and improve data discovery and curation without impacting storage performance.

In tandem with the AFX announcement, NetApp expanded its data services portfolio with the new AI Data Engine designed to integrate its own software with the NVIDIA AI Enterprise suite and run on NVIDIA GPU-accelerated compute nodes. Targeting the fast-evolving AI space, AIDE covers three major areas of data services: semantic search, discovery and classification, and governance. Backed with the new global metadata engine providing a searchable catalog of enterprise data, AIDE also adds the following components: Data Sync to eliminate redundant copies and automate data change detection and synchronization, Data Guardrails for monitoring and classifying data and

enforcing security policies, and Data Curator for data discovery and transformation, including vectorization. AIDE leverages the NVIDIA AI Data Platform reference design featuring NVIDIA-accelerated computing and AI Enterprise software, including NVIDIA Inference Microservices (NIM).

NetApp AIDE can run natively within an AFX cluster on the DX50 compute nodes, and the hardware road map calls for the integration of NVIDIA RTX PRO servers with Blackwell Server Edition GPUs. NetApp's Keystone will offer AI as a service using AFX and AIDE. NetApp's AI Data Engine is also designed to integrate with major AI public cloud platforms, including AWS Bedrock, Google Vertex AI, and Azure AI, as well as with specialized AI and sovereign clouds.

Also, at NetApp Insight 2025:

- NetApp previewed the public availability of an Object REST API to enable direct access to Azure NetApp Files and seamless integration with Azure services to train AI models and build applications on top of existing data sets without the need to move or replicate data.
- The addition of iSCSI block protocol support with Google Cloud NetApp Volumes facilitates unified storage across NAS (NFS/SMB) and SAN environments and enables enterprises to run a variety of applications on Google Cloud without changing their existing storage architecture.
- Google Cloud NetApp Volumes can now serve as a native data source for Google's Gemini Enterprise AI applications, allowing organizations to build custom AI agents that leverage enterprise data directly, without additional coding or complex retrieval-augmented generation (RAG) pipelines, to improve the accuracy and reliability of AI-driven outcomes.
- NetApp extended FlexCache and SnapMirror support to Azure NetApp Files and Google Cloud NetApp Volumes for unified, low-latency access; efficient data replication; and hybrid and multicloud data mobility. Support was already available for Amazon FSx for NetApp ONTAP.
- NetApp added significant enhancements to its cyber protection and cyber-resilience portfolio, embedding AI more deeply into storage and recovery. The key announcement was its Ransomware Resilience service that leverages AI for data breach detection and clean room recovery. The focus is on early detection — a key to minimizing damage — which the company touts as near-100% accuracy. NetApp also introduced secure "clean rooms" for safe postattack recovery. These environments use AI to detect compromised files, identify the last safe data point, and guide users through a verified restore process to minimize the risk of reinfection during recovery.

IDC'S POINT OF VIEW

The AI, cloud, and cyber protection offerings that NetApp launched at its INSIGHT 2025 event represent the latest steps in an ongoing strategy to strengthen the company's position in the ever-evolving market for enterprise storage, where vendors now face the dual challenge of addressing the requirements of traditional business applications and the growing demands for AI-ready data infrastructure that can span hybrid multicloud environments.

NetApp's historical strength rested with NAS file server appliances, but the company has done as much, if not more, than competitors to adapt to the changing conditions in the external enterprise storage market. When public cloud providers threatened traditional storage vendors, NetApp built a software-defined version of ONTAP to run in the cloud and forged partnerships to make available ONTAP as a native, first-party cloud service in the three major public clouds. Support for file, block, and object protocols enables NetApp to promote ONTAP as a unified storage infrastructure for hybrid cloud/multicloud environments, with a central control plane through the NetApp Console (formerly BlueXP, introduced in 2022). ONTAP's flexibility is critical at a time when IDC research shows 70% of enterprises now take a hybrid cloud/multicloud approach to data storage infrastructure. The introduction of lower-cost quad-level cell (QLC) C-Series all-flash arrays, dedicated ASA SAN storage systems, and flash-based StorageGRID object storage helped vault NetApp's 2024 external storage systems revenue by 11% — the company's only double-digit growth since 2018 — after year-over-year declines in 2019, 2020, and 2022 and roughly flat performance in 2023. The new AFX and AIDE additions will better position NetApp in the expanding and increasingly competitive market for AI-ready data storage infrastructure while offering the flexibility of service-based options through Keystone.

AFX is an especially important product addition for NetApp as more storage vendors pursue disaggregated storage architectures to address the performance and capacity requirements of AI workloads. NetApp's convergence of storage controllers and enclosures with NVIDIA-based accelerated compute nodes in the same hardware rack presents a potentially compelling alternative for enterprises struggling to simplify and consolidate their AI data storage infrastructure. NetApp elected to use ONTAP for parallel file storage, through its support for parallel NFS (pNFS), rather than buy or build a parallel file system. The NVIDIA DGX SuperPOD certification should give customers a degree of comfort that an AFX cluster can handle high-performance demands. With a pNFS-based system, enterprises won't need to maintain special client software, as they do with traditional parallel file systems. Plus, they gain access to NetApp ONTAP-based storage capabilities, such as high availability, inline data compression and deduplication, policy-based security, and AI-driven ransomware protection and predictive analytics. NetApp will

also continue to offer the open source BeeGFS on its E-Series hardware for customers interested in deploying a more traditional parallel file system.

The introduction of the AI Data Engine is another significant move for NetApp for several reasons. First, it deepens NetApp's partnership with NVIDIA by integrating not only the GPU specialist's hardware acceleration technologies but also its AI Enterprise software. The tighter integration with the NVIDIA stack will facilitate the optimal use of NVIDIA's acceleration capabilities in NetApp environments. Second, the integration of NVIDIA capabilities with NetApp's ONTAP data management software opens the door to seamless implementation of new AI tools for NetApp users. Third, the broad availability of AIDE across on premises, public cloud, and specialized cloud sites is consistent with NetApp's long-term strategy of delivering common functionality across all customer deployment environments and enables greater integration of corporate resources into AI projects. And, finally, AIDE solidifies NetApp's overall AI strategy.

To date, the rapid growth in spending on server infrastructure, driven by demand for accelerated computing, has not carried over to the external enterprise storage systems market. However, IDC expects steady demand for storage to continue as the AI market focus shifts from model training to inferencing, RAG, and agentic applications, driving demand for tools that can prepare and use internal corporate data in AI-based processes. Although other storage vendors are rushing to develop similar offerings, and the competition will intensify for AI-ready data infrastructure, NetApp should hold a strong position given its well-established customer base, solid ecosystem of partners, and extensive portfolio targeting AI workloads.

NetApp also holds an advantage with its hybrid cloud portfolio. IDC recognizes NetApp as the only storage provider offering first-party integrations across all major cloud providers. The NetApp ONTAP platform remains one of the most widely adopted and technically mature storage technologies, with value for developers seeking consistency and control across environments. ONTAP's unified namespace capabilities, FlexCache and SnapMirror extensions, and expanded object and block interfaces show how NetApp is steadily building a practical framework for hybrid cloud and multicloud data management. These capabilities allow organizations to use the same data sets across on premises and multiple cloud platforms for analytics, AI training, and application development and minimize redundant data copies. The result is a simpler operational model that can reduce both data movement and infrastructure cost while preserving data integrity across locations.

Considering the latest announcements in the context of the yearslong efforts to adapt the NetApp product portfolio to address the continually evolving requirements of enterprise storage buyers shows a company that is willing to take bold steps — and IDC Enterprise

Storage Systems Tracker data signals that some initiatives are starting to pay off. NetApp was one of only two major vendors with double-digit growth in the external enterprise storage market in 2024, when the company was placed third behind Dell Technologies and Huawei. NetApp jumped ahead of Huawei into second place during the first half of 2025, when its all-flash array revenue for SAN and NAS systems (excluding hyperconverged infrastructure appliances, which it does not sell) overtook Dell for the top position in the market. The impact of the new AFX and AI Data Engine remains to be seen, but NetApp has taken the necessary steps to stay in the thick of the intensely competitive market for AI-ready data infrastructure.

Subscriptions Covered:

[Cloud Data Logistics and Protection](#), [Enterprise Workloads Infrastructure Stacks and Deployments](#), [Public and Dedicated Cloud IaaS](#), [Storage Systems](#)

Please contact the IDC Hotline at 800.343.4952, ext.7988 (or +1.508.988.7988) or sales@idc.com for information on applying the price of this document toward the purchase of an IDC or Industry Insights service or for information on additional copies or Web rights. Visit us on the Web at www.idc.com. To view a list of IDC offices worldwide, visit www.idc.com/offices. Copyright 2025 IDC. Reproduction is forbidden unless authorized. All rights reserved.