

**451 Research Market
Insight Report Reprint**

NetApp AFX disaggregated storage takes on data- intensive AI workloads

November 25, 2025

by **Henry Baltazar**

At the NetApp Insight 2025 conference, the vendor launched the NetApp AFX system, which is designed for data-intensive AI workloads. This follows the recent StorageGRID 12.0 update, which introduced S3 Bucket Branches for versioning as well as a new native caching capability.



This report, licensed to NetApp, developed and as provided by S&P Global Market Intelligence (S&P), was published as part of S&P's syndicated market insight subscription service. It shall be owned in its entirety by S&P. This report is solely intended for use by the recipient and may not be reproduced or re-posted, in whole or in part, by the recipient without express permission from S&P.

Introduction

At the NetApp Insight 2025 conference, the vendor launched the NetApp AFX system, which is designed for data-intensive AI workloads. This follows the recent StorageGRID 12.0 update, which introduced S3 Bucket Branches for versioning as well as a new native caching capability.

THE TAKE

NetApp Inc.'s new disaggregated AFX system and AI Data Engine are key additions that should make the vendor's ONTAP storage environment better suited for modern AI enterprise workloads. The vendor's expanding partnerships with players such as NVIDIA Corp., Google Cloud and Microsoft Corp. (Azure) will continue to make its technology available via reference architectures and native AI platforms. All of its efforts to ramp up product development and partnerships are essential, given that competition for storage and data management is getting more intense.

Context

NetApp's 2025 announcements represent a forward-looking extension of its strategy, which unifies its architecture under the NetApp data platform and introduces purpose-built innovations such as AFX and the AI Data Engine. The vendor is using these innovations to drive its evolution from a storage-centric vendor into a broader enabler of intelligent, AI-ready infrastructure. The success of this transition will depend on sustained delivery, ecosystem execution and the ability to communicate value beyond hardware. The recent enhancements to the StorageGRID object storage platform provide native versioning through the new S3 Bucket Branches capability, in addition to new caching capabilities and metadata scalability enhancements to support up to four billion objects per high-performance node.

In its earnings report for Q1 FY 2026 (period ending July 25, 2025), the vendor reported all-flash revenue of \$893 million, which was a 6% year-over-year increase that amounts to an annualized net revenue run rate of \$3.6 billion. NetApp also reported that its first-party and marketplace public cloud storage services revenue grew 33% year over year. Overall the vendor's net revenue of \$1.56 billion was a 1% increase from a year ago.

NetApp AFX disaggregated storage for AI workloads

NetApp AFX is a disaggregated storage system optimized for unstructured data and large-scale AI workloads. Rather than replacing ONTAP, AFX extends it into a new deployment model that supports independent scaling of compute and storage resources. AFX is built entirely on the existing ONTAP codebase, maintaining full API and feature compatibility with other NetApp systems, such as AFF, ASA and FAS. The architecture scales from initial eight-node configurations to as many as 128 controller nodes within the first year, supporting exabyte-scale capacity. It employs a 400-Gb Ethernet fabric with RDMA acceleration to achieve high throughput and low latency without requiring proprietary networking components.

Complementing AFX, NetApp introduced the AI Data Engine, a suite of integrated data services designed to prepare enterprise data for AI workloads. The AI Data Engine extends ONTAP's capabilities beyond storage management into data discovery, governance and transformation. It is built around three core components: Metadata Engine, which provides visibility and classification; Data Guardrails, which automates governance and compliance; and Data Curator, which supports data transformation and vector embedding generation for retrieval-augmented generation workflows.

Platform enhancements and cloud integrations

In addition to the AI-focused releases, NetApp introduced several updates across its portfolio. The upcoming ONTAP 9.18.1 release delivers scalability improvements for ASA systems, enhanced object-storage integration and new encryption capabilities. NetApp also launched the new NetApp Console, replacing the BlueXP brand and providing a unified interface for managing the NetApp Data Platform.

Ecosystem and cloud integration

NetApp continues to rely heavily on partnerships to extend its market reach. A key part of this strategy is deeper alignment with NVIDIA, Cisco Systems Inc., Google Cloud and Microsoft Azure. A new Cisco FlexPod AI reference architecture integrates AFX storage, Cisco servers with NVIDIA GPUs, and the AI Data Engine to provide an enterprise-ready blueprint for AI deployments. In the public cloud, Google Cloud NetApp Volumes and Azure NetApp Files have been elevated to “first-class” data sources for native AI platforms such as Google Vertex AI, Gemini and Azure AI.

StorageGRID adds bucket branching and caching

StorageGRID continues to evolve into a two-tiered architecture that separates high-capacity storage from low-latency cache layers. The outer ring functions as the high-throughput, exabyte-scale data layer, providing a unified namespace across sites. The inner ring acts as a near-GPU caching tier, offering proximity-based acceleration for AI and high-performance computing environments. This design aims to support workflows that require both large data repositories and responsive data access during model training or inference. NetApp describes this model as a foundation for future hybrid AI and HPC deployments.

The core enhancements in Version 12.0 include several key features to provide versioning, improve performance and scalability, and tighten security and access management, including:

- **S3 Bucket Branches:** StorageGRID introduces a new native versioning feature known as S3 Bucket Branches. This allows users to create and clone buckets based on a seed bucket, similar to Git branching. Point-in-time buckets are created without duplicating data and can be configured as read-only or read-write. Each bucket maintains independent role-based access controls, and bucket rollback is supported. This function supports reproducible AI workflows by enabling controlled data versioning and rollback.
- **Native read caching layer:** Version 12.0 adds a caching mechanism that reduces read latency for workloads such as training, media processing and content distribution. Cache performance can be scaled by enabling additional load-balancer nodes. Benchmarks indicate a significant throughput improvement per cache node compared with traditional configurations. Administrators can monitor cache activity through system metrics such as hit rate, latency and cache utilization.
- **Metadata and scalability improvements:** StorageGRID now supports up to four billion objects per high-performance node, doubling the previous limit. The change applies to SG6000 and SG6100 appliance classes, as well as equivalent virtual or containerized deployments. The capacity increase results from schema and compression updates to the metadata database, as well as space-efficiency improvements in Cassandra.
- **Security and access management:** Support for the AWS AssumeRole API has been added, enabling temporary credentials and delegated access control across tenants. In-line session policies can define time-limited or role-specific permissions for S3 bucket operations.
- **Appliance and upgrade enhancements:** Appliance management includes automated drive firmware updates and reduced downtime during upgrades. Administrators can convert existing nodes into data-only configurations to optimize metadata placement and throughput. New hardware options include 24-TB hard drives for the SG5800 and SG6100 series, and 122.8-TB QLC SSDs for the SGF6112 model.

- **Global consistency and reliability:** A new “Strong Global v2” consistency mode introduces quorum-based validation for read, write and delete operations across multiple sites. This feature allows continued operation in the event of a single-site outage within a three-site deployment, improving operational continuity in distributed grids.
- **Event notification via webhooks:** The platform now supports webhooks for bucket event notifications. Users can configure webhook endpoints to trigger on PUT or DELETE events, enabling integration with local or third-party services through HTTPS payloads. The system supports Secure Socket Layer and mutual Transport Layer Security authentication for security.
- **Traffic management expansion:** The limit for tenant traffic classifiers has been increased from 500 to 2,500 policies. This change simplifies traffic isolation and control in large multi-tenant environments such as service providers. Object-size metrics are not recorded for policies beyond the first thousand due to database constraints.

Competition

Storage platforms with disaggregated architectures are moving into the mainstream, with notable rivals — including VAST Data, Pure Storage Inc. with its FlashBlade//EXA and Hewlett Packard Enterprise Co. with its Alletra MP storage systems — using this architecture to better match the high-scalability and low-latency performance requirements of GenAI workloads. NetApp also continues to compete with Dell Technologies Inc., IBM Corp., Hitachi Vantara and several other storage players in the market.

Object storage continues to be a highly competitive space, where NetApp will compete with public cloud storage services from AWS, Azure and GCP, in addition to smaller players such as Wasabi Storage and Seagate Technology Holdings PLC (Lyve). In the on-premises market, there continues to be competition from vendors such as Scality, MinIO and Cloudian — in addition to the object storage platforms of the previously mentioned enterprise players. The new S3 bucket branches, native caching and metadata scalability improvements are all necessary improvements for NetApp’s StorageGRID in this competitive segment.

SWOT Analysis

STRENGTHS NetApp is a top storage player and continues to innovate to keep its lineup competitive.	WEAKNESSES The vendor has not historically had success selling compute offerings in the past, unlike rivals that have large server businesses.
OPPORTUNITIES NetApp AFX and the new AI Data Engine are designed to help enterprise customers manage and prepare data for AI workloads. The StorageGRID version 12 update improves performance and scalability, and adds new innovations in the form of S3 Bucket Branches.	THREATS The vendor is facing strong competition not only against traditional infrastructure players, but also from public cloud storage services.

CONTACTS

Americas: +1 800 447 2273

Japan: +81 3 6262 1887

Asia-Pacific: +60 4 291 3600

Europe, Middle East, Africa: +44 (0) 134 432 8300

www.spglobal.com/marketintelligence

www.spglobal.com/en/enterprise/about/contact-us.html

Copyright © 2025 by S&P Global Market Intelligence, a division of S&P Global Inc. All rights reserved.

These materials have been prepared solely for information purposes based upon information generally available to the public and from sources believed to be reliable. No content (including index data, ratings, credit-related analyses and data, research, model, software or other application or output therefrom) or any part thereof (Content) may be modified, reverse engineered, reproduced or distributed in any form by any means, or stored in a database or retrieval system, without the prior written permission of S&P Global Market Intelligence or its affiliates (collectively S&P Global). The Content shall not be used for any unlawful or unauthorized purposes. S&P Global and any third-party providers (collectively S&P Global Parties) do not guarantee the accuracy, completeness, timeliness or availability of the Content. S&P Global Parties are not responsible for any errors or omissions, regardless of the cause, for the results obtained from the use of the Content. THE CONTENT IS PROVIDED ON "AS IS" BASIS. S&P GLOBAL PARTIES DISCLAIM ANY AND ALL EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR USE, FREEDOM FROM BUGS, SOFTWARE ERRORS OR DEFECTS, THAT THE CONTENT'S FUNCTIONING WILL BE UNINTERRUPTED OR THAT THE CONTENT WILL OPERATE WITH ANY SOFTWARE OR HARDWARE CONFIGURATION. In no event shall S&P Global Parties be liable to any party for any direct, indirect, incidental, exemplary, compensatory, punitive, special or consequential damages, costs, expenses, legal fees, or losses (including, without limitation, lost income or lost profits and opportunity costs or losses caused by negligence) in connection with any use of the Content even if advised of the possibility of such damages.

S&P Global Market Intelligence's opinions, quotes and credit-related and other analyses are statements of opinion as of the date they are expressed and not statements of fact or recommendations to purchase, hold, or sell any securities or to make any investment decisions, and do not address the suitability of any security. S&P Global Market Intelligence may provide index data. Direct investment in an index is not possible. Exposure to an asset class represented by an index is available through investable instruments based on that index. S&P Global Market Intelligence assumes no obligation to update the Content following publication in any form or format. The Content should not be relied on and is not a substitute for the skill, judgment and experience of the user, its management, employees, advisors and/or clients when making investment and other business decisions. S&P Global keeps certain activities of its divisions separate from each other to preserve the independence and objectivity of their respective activities. As a result, certain divisions of S&P Global may have information that is not available to other S&P Global divisions. S&P Global has established policies and procedures to maintain the confidentiality of certain nonpublic information received in connection with each analytical process.

S&P Global may receive compensation for its ratings and certain analyses, normally from issuers or underwriters of securities or from obligors. S&P Global reserves the right to disseminate its opinions and analyses. S&P Global's public ratings and analyses are made available on its websites, www.standardandpoors.com (free of charge) and www.ratingsdirect.com (subscription), and may be distributed through other means, including via S&P Global publications and third-party redistributors. Additional information about our ratings fees is available at www.standardandpoors.com/usratingsfees.