

NETAPP AFX

Intelligent data infrastructure built for the
AI-powered enterprise



Accelerate AI outcomes with high-performance, ultrascalable, disaggregated storage powered by enterprise-proven NetApp ONTAP

AI is now an enterprise workload

AI is no longer experimental. It is now an enterprise workload, standing alongside virtualization, databases, and other mission-critical systems. Like these workloads, AI demands enterprise-grade capabilities such as data protection, quality of service, resilience, and high availability. But unlike traditional workloads, AI requires agile scalability, ultra-high throughput, performance to support dynamic, and data-intensive pipelines that power real-time inference and model training. Agentic AI only exacerbates the issue, requiring deep integration and connectivity with enterprise systems to automate core functions and deliver ROI.

Enterprises face critical infrastructure barriers to operationalizing AI:

- Legacy storage architectures can't scale to meet the high-performance demands of modern AI workflows.
- Fragmented point solutions introduce silos, lack

enterprise-grade security and resilience, and increase operational complexity.

- Hybrid cloud and data mobility challenges disrupt AI pipelines, making it harder to move, manage, and curate data across environments.

NetApp AFX is disaggregated storage built for the AI-powered enterprise

NetApp® AFX combines extreme performance and scale with the reliability of enterprise-proven NetApp ONTAP® software. AFX is built on ONTAP, benefiting from over three decades of world-class software and hardware engineering with enterprise-proven data management and security. Because it is ONTAP, AFX integrates seamlessly into your enterprise data estate and delivers granular, policy-based security so that AI accesses only the data you intend.

A revolutionary architecture delivered on enterprise-proven ONTAP

Key to the high performance and scalability of AFX is its innovative disaggregated architecture. The compute layer that handles data management and serves I/O is decoupled from the capacity layer of high-performance NVMe Flash storage. This design allows customers to scale performance and capacity independently. Add storage controllers to boost performance and storage enclosures to expand capacity, all with linear gains

This flexibility means that customers can optimize their infrastructure to meet their unique workload requirements without overprovisioning. Built for non-stop operations means upgrades and expansions are simple sparing you the outages and downtime required by parallel file system solutions. And because AFX is built on ONTAP, it delivers the same best-in-class security and ease of use that enterprise customers demand.

Extreme performance, standard protocols, seamless integration

NetApp AFX delivers all the performance benefits of parallel file systems and niche AI storage solutions, but on an enterprise-grade platform that is simple, secure, and fully integrated.

Unlike solutions that require proprietary file system clients, AFX uses standard file and object protocols, including parallel NFS (pNFS) for extreme performance and AWS S3-compatible object storage for flexibility. This integrated platform means that all of your applications can use AFX without installing custom clients that introduce instability, security risks, or operational complexity.

With AFX, you get:

- **Enterprise-grade simplicity.** No proprietary clients, no fragile integrations, just standard protocols and a single pool of exascale capable storage.
- **Performance without trade-offs.** Parallel NFS delivers the throughput that AI workloads demand while maintaining compatibility and eliminating complexity.
- **Hybrid cloud readiness.** Only NetApp delivers native ONTAP integrations across every major cloud giving you ultraefficient data access and mobility across on-premises and cloud environments with NetApp SnapMirror® replication and intelligent caching via FlexCache®.

KEY BENEFITS

- **Validated for NVIDIA DGX SuperPOD. Built on ONTAP.** Get the speed, scale, and resilience needed for enterprise AI - without compromise.
- **AI performance and enterprise proven capabilities** empowers your team to operationalize AI without retraining
- **Storage simplicity plus enterprise resilience** with the AFX architecture that delivers space efficiency and high performance.
- **Find and curate precise data sets** with an always-up-to-date index that doesn't impact primary workload performance.
- **The most secure storage on the planet.** Safeguard sensitive data with secure multi-tenancy while optimizing resources with granular QoS for compliant and responsible AI.

Operationalize AI with NetApp AI Data Engine

Accelerate AI pipelines by consolidating fragmented tools into a unified, ONTAP-integrated solution with real-time metadata, inline vectorization, and semantic search. NetApp AI Data Engine delivers a global, current view of your data estate, enabling fast, precise curation, deep integration with leading AI platforms, and built-in governance for secure, compliant AI operations.

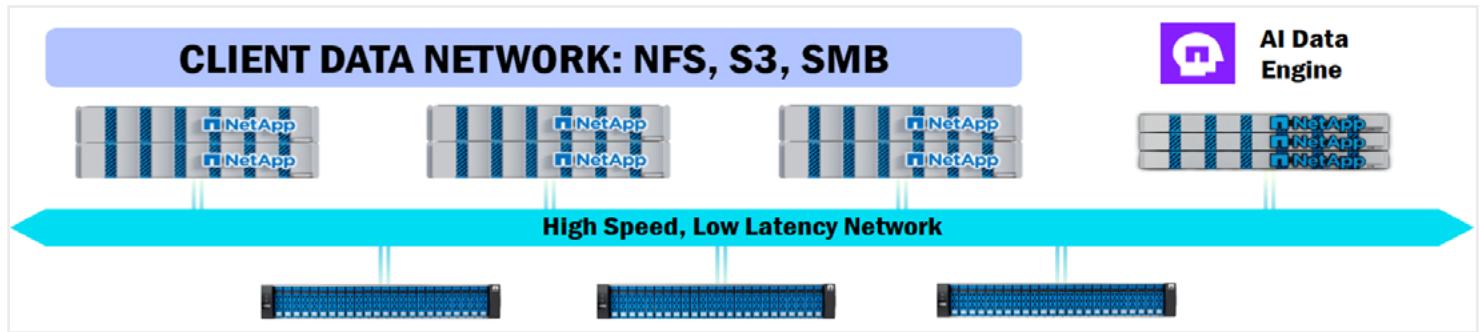


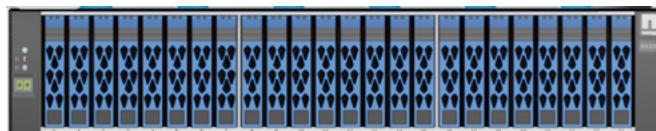
Fig. 1: NetApp AFX deployment consists of AFX 1K storage controllers, NX224 NVMe storage enclosures, and optional DX50 data compute nodes.

AFX components



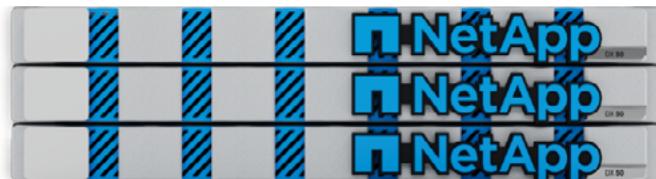
AFX 1K storage controller

Built on NetApp engineered hardware, the AFX 1K storage controller delivers extreme performance in a resilient and easily serviced form factor.



NX224 NVMe enclosure

The NX224 NVMe enclosure provides ultrafast, low-latency storage for AI datasets. Customers can add enclosures to scale capacity seamlessly as data volumes grow.



DX50 data compute nodes

AFX also introduces the DX50, a GPU-powered data compute node in a compact 1U form factor. The DX50 powers the AFX advanced metadata engine, which continuously indexes data and enables data engineers to discover and curate datasets by using familiar SQL queries. Because this engine runs on dedicated compute nodes, it never impacts the performance of your primary AI workloads.

Kickstart AI with NetApp Keystone STaaS

Keystone STaaS for Enterprise AI removes barriers to AI adoption by eliminating upfront capital costs and simplifying deployment. Delivered as a single, consumption-based subscription and powered by NetApp AFX and AIDE, it enables AI-optimized performance and independent capacity scaling on demand. Predictable billing allows usage adjustments and burst capacity at no extra cost, while pay-as-you-go pricing aligns expenses with workload demand. Integrated management via NetApp Console and DII automate operations (so your staff can focus on business innovation), and built-in cyber resilience provides advanced threat detection and ransomware protection to empower enterprises to operationalize AI quickly, securely, and cost-effectively.

AFX technical specifications

AFX cluster	
Maximum storage controllers	128
Maximum storage enclosures	52
Maximum effective capacity	1+ EB
Maximum data compute nodes	10
Supported storage protocols	pNFS, NFS, SMB, S3, NFS/RDMA

AFX 1K storage controller	
Controller form factor	2U
PCIe expansion slots	11 (9 + NVRAM)
X50131B — 2p, 100G/200G/400G Ethernet	4
X50130B — X50130B — 2p, 40G/100G Ethernet controller 100GbE ports (40Gbe autoranging)	5
X50133A — 4p, 10G/25G Ethernet controller 10GBASE-T (1Gbe autoranging)	5
OS version	ONTAP 9.17.1 or later
Power consumption (median or typical)	Voltage - 200 (200 to 240) Amps - 6.48 Typical

DX50 data compute node	
Form factor	1U
CPU	AMD Genoa 9554P
Sockets	1
Total cores	64
Memory	1TB
I/O	4-ports, 100GbE
GPU	1x NVIDIA L4
Internal storage	2xU.2 15TB
Typical power consumption	110 Voltage: Typical - 4.31 AMPS 220 Voltage: Typical - 9.10 AMPS

NX224 storage enclosure

Maximum drives	24
Drive form factor	2.5" small form factor
Form factor	2U
Drive type supported	7.6 TB, 15.3 TB, 30.7 TB, 61.4 TB
OS version	ONTAP 9.17.1 or later
Typical power consumption	TBD

AFX software

High availability	<ul style="list-style-type: none"> Disaggregated active-active controller designed for 6 9s of availability Nondisruptive maintenance, upgrade, and scale-out clustering
Storage efficiency	<ul style="list-style-type: none"> Inline data compression, deduplication, and compaction Space-efficient file and volume cloning Maximize your flash ROI with data protection overhead as low as 10%
Data management	<ul style="list-style-type: none"> Intuitive onboard GUI, REST APIs, and automation integration Proactive data balancing across cluster AI-informed predictive analytics and corrective action Quality-of-service workload control Easy provisioning and data management from market-leading host operating systems, hypervisors, and application software Asymmetrically stripe large files across FlexGroup API support for NAS volumes sharing their data via S3
Scalable NAS	<ul style="list-style-type: none"> Exascale capable single namespace management with local and remote caching
Data protection	<ul style="list-style-type: none"> Application-consistent NetApp Snapshot™ copies and restore Integrated remote backup and disaster recovery Synchronous zero-data-loss replication Tamper-proof Snapshot copies
Security and compliance	<ul style="list-style-type: none"> AI-powered Autonomous Ransomware Protection Multifactor administrator access Secure multitenant shared storage In-flight and data-at-rest encryption Regulatory-compliant data retention Multiadmin verification before executing sensitive commands
Cloud integration	<ul style="list-style-type: none"> Seamlessly back up, replicate, and cache data to private and public clouds Native data mobility to major public cloud services

* These specifications represent reasonable customer configurations. The AFX architectural design can scale more than an order of magnitude above the noted limits. Upon initial release of the AFX 1K certain smaller limits may be imposed; please check with NetApp for details. Full enablement of the limits in the referenced table is planned for a future software release of ONTAP, included at no additional cost for customers under a support contract. Statements by NetApp about unreleased offerings and future plans are for informational purposes only, are subject to change without notice, and should not be relied upon for purchasing or other decisions. Such statements do not constitute a commitment, obligation, guarantee, or warranty of any kind by NetApp, including about availability, functionality, pricing, or timing.



Contact Us

About NetApp

NetApp is the intelligent data infrastructure company, combining unified data storage, integrated data services, and CloudOps solutions to turn a world of disruption into opportunity for every customer. NetApp creates silo-free infrastructure, harnessing observability and AI to enable the industry's best data management. As the only enterprise-grade storage service natively embedded in the world's biggest clouds, our data storage delivers seamless flexibility. In addition, our data services create a data advantage through superior cyber resilience, governance, and application agility. Our CloudOps solutions provide continuous optimization of performance and efficiency through observability and AI. No matter the data type, workload, or environment, with NetApp you can transform your data infrastructure to realize your business possibilities. www.netapp.com

