

NetApp All Flash FAS (AFF)

Overview

The NetApp fabric-attached storage (FAS) series with SSDs only are offered as the All Flash FAS (AFF) systems and have several different models. NetApp refers to fabric-attached storage systems as storage connected to hosts via any network. The FAS series is part of NetApp's strategy to offer a network storage system that is capable of supporting both NAS and storage-area network (SAN) accessibility from a single storage pool. All Flash FAS (AFF) only supports SSDs and have embedded software changes to optimize use of flash.

The AFF family spans from entry level to high end, with AFF A series systems including the AFF A150, A250, A400, A800, and A900. NetApp additionally offers C models for capacity oriented workloads, and designated ASA models for SAN only.

The embedded operating software is called ONTAP and support scale-out clustering and advanced feature capabilities. NetApp ONTAP has scale out capability to allow systems to be configured to match the environment needs and to increase in both capacity and performance as required. The scale out capability with ONTAP allows scaling to 24 node nodes (the A250 can scale to 8 nodes) along with the capacity scaling for NAS and 12 nodes for block mode. A dual node NetApp system is called a High Availability configuration or HA pair. A cloud version of ONTAP is also available.

Data reduction in the form of deduplication, compression, and compaction are supported on NetApp systems for primary data. The data reduction is performed inline and NetApp provides a 4:1 reduction guarantee. Compaction is filling the 4K allocation chunks with smaller blocks, which can improve storage utilization, especially with small files.

The FabricPool feature tiers snapshot data not in the active filesystem and inactive blocks from a primary storage system and SnapMirror or SnapVault data from a secondary storage system to either Amazon S3 or StorageGRID WebScale.

Highlights

- Unified storage – block, file & object – scale out
- All flash version of FAS
 - NVMe and NVMe over Fabric
- Extremely large set of advanced features
- SAN-only models for streamlined block support.
- ASIC acceleration in A400 systems.
- ONTAP software with Intel-based hardware
- Inline Data Reduction – compression deduplication, compaction
- MetroCluster for stretched cluster
- Snaplock for immutability
- Encrypting software and internal keys
- OnCommand management software
- Large size and number of files with FlexGroups
- FabricPool tiering to cloud & object storage

Usage

The NetApp systems are pervasively used in IT as primary storage, secondary storage, and even for archiving. Both block and file support are in usage but predominantly the NetApp systems are used as NAS for file storage. The systems fit typical IT usages and are found in very wide applicability across vertical markets. The FlexGroup feature allow for use in media & entertainment, oil & gas, EDA, and healthcare usage. The feature also increases parallelism to improve performance with large numbers of files in multi-node configurations.

- Characteristics
 - Performance – As standard NAS systems, the NetApp AFF systems have excellent performance for file storage and retrieval for mixed types of files. There are SPEC benchmarks published with details of performance. The block usage also have excellent performance. NetApp quotes performance as 11.4M IOPs and 300 GB/s bandwidth for the AFF A800 model in a cluster.
 - Availability – The dual node systems are active-active for availability and the scaling with shared access to the storage arrays provides a high availability solution. Both dual parity (RAID-DP) and triple parity RAID-TP protection is available.
 - Scale out configurations in cluster mode
 - Performance and capacity can be scaled together with the addition of more nodes to the cluster. Scaling with performance and capacity allows customers to maintain performance as capacity demands are being met.
 - Global namespace across all nodes in the cluster simplifies access to files or LUNs contained in the different nodes. Operational tasks and user access is greatly enhanced without having to manage multiple mounts or paths to resolve to the data for access.
 - Non-disruptive operations are possible in cluster mode that either could not be done otherwise or only with manual operations.
 - Load balancing can be done with redistributing active data across different nodes.
 - Capacity balancing can be accomplished transparently to manage new nodes with potentially new technology being added to the system.
 - In addition to adding new nodes to the cluster, nodes may be retired from service with the data transparently migrated to other nodes. Technology updating without service interruption or administrative effort in migration is a key enablement for cluster mode operations.
 - Non-disruptive servicing is available as with the high availability configurations.

- Existing added value features such as Snapshots work transparently in the cluster environment.
- Existing interface protocols such as Fibre Channel and iSCSI or NAS protocols operate without change in the cluster environment.
- Object storage is over HTTP/S with S3 operators.
- A single point of management is used across the cluster, eliminating the individual management tasks for non-cluster environments.
- Replication for BC/DR – There are multiple replication features for NetApp including incremental Snapshots, asynchronous, and synchronous remote replication providing many choices for customers in meeting BC/DR needs. Relaxed Synchronous allows for less performance impacting replication.
- Notable – The wealth of advanced features is greater than any other individual storage system. These features extend the value of the system providing additional capabilities for companies and their workflows.
- Applications
 - The NetApp AFF storage systems are targeted at file-based or object applications (unstructured data) and block storage usages. As unified storage, all types of usages can be supported at the same time.
 - With file, object, and block storage support, the NetApp AFF can be used as a consolidation solution across different requirements.
- System environments
 - With the broad range of file protocols including native CIFS/SMB and NFS for file and Fibre Channel and iSCSI for block storage, the NetApp systems can install in almost any environment.
 - HDFS support for Hadoop environments
- Deployment and Administration
 - Noted for the ease of use, the NetApp systems can be installed by users with moderate networking and storage skills.
 - NetApp provides resource management software as well as device and feature managers. A CLI is also available for low-level management.
 - With Active IQ, telemetry data is sent to a cloud-based application that performs analytics on the data. Active IQ is an update of the original Autosupport function.

Key Capabilities

Hardware Architecture

The AFF systems are all flash storage systems with dual active-active controllers for high availability. The AFF appliances run NetApp ONTAP software and host connectivity is provided via Fibre Channel or Ethernet connectivity. A key hardware element in AFF systems is NetApp's

use of NVRAM. NVRAM is used alongside the WAFL file system to allow writes to be buffered and pushed out in batches.




Figure 1: NetApp FAS/AFF Architectural Components

Scale Out

AFF systems are scale-out systems capable of adding additional systems to a cluster with a shared namespace. The systems scale in increments of HA pairs, with different limitations for block and file usage. Block systems can scale out to a total of 6 HA pair systems, or 12 nodes, while file storage can scale to 12 HA pair systems, or 24 total nodes.

ASIC Acceleration

The A400 model features a Pensando ASIC chip to offload some of the data processing workload from the main controller CPUs. The ASIC chip is used to aid in data reduction calculations including compression, decompression, and calculating checksums for deduplication.

ONTAP

The ONTAP operating system has evolved over time to become the basis for NetApp's product portfolio. ONTAP began as the name for the embedded real-time OS that NetApp used in their products. As the ONTAP OS progressed, NetApp began to treat their OS as a platform, upon

which additional products and features are delivered. ONTAP is used in NetApp AFF, FAS, and ASA systems, as well as the basis for its cloud offerings such as Cloud Volumes ONTAP and AWS FSx for ONTAP.

There are three primary components of ONTAP, the OS execution environment, the filesystem, and the data protection features. NetApp's filesystem is referred to as WAFL, and the data protection is implemented via a software RAID mechanism. An additional RAID level was made available as part of ONTAP 6.5 and is referred to by NetApp as RAID-DP, which is RAID dual parity, a specific implementation often known as RAID-6. ONTAP 9 added support for Triple Parity RAID. NetApp's ONTAP based systems may be used for a variety of purposes, including providing file access to hosts and providing virtualized block data to hosts via both FC and IP networks.

WAFL Filesystem

The NetApp Write Anywhere Filesystem Layout (WAFL) is one of the features that are a key differentiator for NetApp. The use of WAFL provided administrative simplicity by minimizing the number and complexity of settings. NetApp uses the log-structured file system WAFL for ONTAP's internal organization and allocation of data.

As originally designed and implemented, WAFL has many similarities to the Unix BSD filesystem FFS or Fast Filesystem.¹ A detailed overview of WAFL's block map architecture is provided below.




Figure 2: NetApp ONTAP Architecture (Source NetApp)

One key feature of WAFL is the manner in which it stores information about the files, or the metadata. WAFL stores metadata in regular files, instead of as part of the inode structure as many other file systems do. This provides greater flexibility as data structures may be changed more easily without requiring complex interface changes or data migrations upon updates. By maintaining metadata in files, the WAFL filesystem is able to write the metadata anywhere a file can be written on disk.




Figure 3: NetApp ONTAP Architecture (Source NetApp)

Snapshots

NetApp ONTAP supports snapshots using a redirect on write implementation. As implemented by NetApp, snapshots are read-only references to the original data. These references do not contain duplicate data and rely on the original data updating only new information. Snapshots are created by creating a copy of the pointers to the data, not the actual data itself. As such, snapshots are highly space efficient, and fast to create. Typically, a point-in-time snapshot copy is created in less than 5 seconds.

MetroCluster

MetroCluster is an integrated high-availability solution that extends failover capability from within a data center offsite and replicates data from the primary site to the remote site allowing for a combination of failover and data replication. As a stretched cluster implementation, MetroCluster supports business continuity at distances up to 700KM. In addition to the failover mechanisms for block and file support, MetroCluster supports workload mobility, allowing virtual machines to be moved across geographic distances without having to move the data. Additionally, non-disruptive operations between data centers are supported as part of the failover mechanism. MetroCluster can work over Fibre Channel or IP (Ethernet) links.

SnapMirror

SnapMirror is a remote copy solution for file systems or LUNs provided by NetApp systems. A file system or LUN can have SnapMirror established which will make a copy of the file system (and all metadata) or LUN to another appliance over a LAN or WAN. Once the copy is made, the mirror copy is kept current by updating with only the new or changed data blocks. Like the base Snapshot feature, access to the original file system or LUN can continue because of the WAFL

pointer (inode) indirection being done. The SnapMirror feature allows for control of the frequency of updating the mirror copy. The interval setting can be as low as one minute. By doing an asynchronous copy, the updating of the mirror copy with changed or new blocks will not affect performance: applications using the source volumes can continue accessing the original file system or LUN. SnapMirror also supports synchronous mirroring beginning with ONTAP version 9.5. Synchronous replication is available in two modes, full synchronous and relaxed synchronous.

SnapMirror is useful for remote copies that may be required for a disaster recovery for data migration to a remote site where the data would only be used in read-only mode (re-synchronization back to the original copy is not an option), or for a remote backup solution. NetApp inherently supports consistency groups for remote copy of multiple LUNs such as would be used in transaction processing databases. The volume group logical construct, which groups multiple LUNs as defined by the user in a Qtree, will maintain write order for the remote copy.

Object Storage

NetApp ONTAP supports object storage access with HTTP and HTTPS over Ethernet. This is an enterprise object storage system like StorageGRID. ONTAP supports a limited set of S3 operators so applications that use more advanced capabilities should use StorageGRID.

Data Reduction

NetApp AFF supports multiple techniques for data reduction including compression, deduplication, and compaction. Data reduction is done inline as it is transferred. Data is deduplicated and compressed by software while it is in memory during the write operation and then compacted and written to devices. The system has an additional storage efficiency mechanism to distinguish between hot and cold data. For cold data, an additional pass is done in the background with a 32K compression group for greater efficiency.

The inline compaction algorithm will attempt to place multiple logical blocks from the same volume in a 4K block to utilize space in the allocation chunks. Small files and databases with small I/O sizes will benefit from the greater capacity utilization.

FabricPool

The FabricPool feature will tier data from NetApp ONTAP storage systems to cloud, specifically Amazon S3, Microsoft Azure, IBM Cloud Object Storage, NetApp StorageGRID object storage system, or another NetApp ONTAP system with S3 support. There is an additional NetApp license charge for use of Amazon S3 and Azure as a FabricPool target. Primary storage can have snapshot data where the data blocks are not part of the active file system tiered (moved) to the cloud target. NetApp estimates that 10% of primary data are snapshots that could be tiered with FabricPool. Secondary storage, which is the data

protection target for SnapMirror or SnapVault will have the data blocks tiered to cloud storage. NetApp estimates that 90% of secondary storage is eligible to be tiered with FabricPool.

SnapLock

The SnapLock feature creates volumes presented as file system shares that are archive targets with file based policies for retention management. The SnapLock volumes are stored as Write Once Read Many (WORM), and can be configured in a compliance mode or enterprise mode.

- SnapLock Compliance
 - Files stored in a SnapLock Compliance volume with the WORM state cannot be deleted or modified until the retention period has expired.
 - The SnapLock Compliance volume cannot be deleted until all files on the volume have their retention periods expired.
- SnapLock Enterprise
 - An administrator can delete SnapLock Enterprise volumes
 - The user or administrator cannot delete or modify files until the retention period expires.

Archived files can have retention settings inherited from a default setting for the entire volume/file system. The retention settings can also be defaulted when the SnapLock volume is created or can be set on a file by file basis by overriding the Last Access Time on the file.

NetApp RAID-DP

To protect data from multiple drive failures that become more likely with larger capacity disk drives, NetApp has developed an implementation which NetApp calls RAID-DP, similar to what other vendors call RAID-6. RAID-DP stands for Dual Parity and is a method of storing blocks with more than one redundancy element. With RAID-6, the striped parity from RAID-5 is utilized, and then an additional parity is calculated and placed on a different disk, effectively providing parity on two separate drives. Each disk in a RAID-6 stripe has two parity regions calculated separately. The advantage is that the array group can tolerate even a two-disk failure, decreasing the risk of failure while rebuilding data onto a replacement for a single failure. With some implementations of RAID-6, an additional performance penalty is incurred while calculating the additional parity information.




Figure 4: RAID-DP Architecture

NetApp's RAID-DP is similar to both RAID-6 and RAID-4, with the minimum number of drives in the RAID group being four disks. For RAID-DP, NetApp still maintains RAID-4 with a separate disk, with a horizontal stripe across all drives with parity on a dedicated drive. In addition, another drive is added to the RAID set, which maintains additional information, including both data and parity.

Active-IQ

Active IQ is a remotely hosted, Web-based service that proactively monitors and identifies storage infrastructure issues to avoid interruptions. The application provides tools to model and optimize the storage infrastructure and includes a mobile application for predictive and personalized support. The original Autosupport for ONTAP was enhanced and renamed as Active IQ with cloud-based analytics on telemetry data from systems.

Extensions to Active IQ have included risk prediction for systems and a recommendation engine.

BlueXP

BlueXP is a centralized SaaS based control plane for deploying and managing NetApp ONTAP based storage both on-premises, such as AFF, and NetApp's ONTAP cloud offerings. Base BlueXP functionality is included with Netapp ONTAP offerings, but also includes functionality for additional management features such as backup, tiering, and other capabilities. The Futurum Group covers NetApp BlueXP separately under Multi-Cloud Data Management research.

OnCommand Unified Manager

Enables centralized management of physical and virtual storage environments through a unified interface. Used for integration of operations, provisioning, and protection, OnCommand Unified Manager allows a unified policy infrastructure to be created. A dashboard with interface

choices and up to nine panels may be created in the manager. Unified Manager also serves as a single configuration repository for reporting, events, and audit trail logs.

Significant Announcements

- 2023
 - NetApp announces AFF C Models and ASA systems
 - C Models use QLC devices for denser storage and capacity optimization
 - ASA provide ONTAP systems configured as SAN only systems
 - ASA C models released to provide QLC based block storage

Futurum Group EvaluScale – Data Protection Software

The Futurum Group product review methodology "EvaluScale" assesses each product within a specific technology area. The evaluation of each product is based on its capabilities, with capabilities for each technology segment grouped into distinct categories. For the NAS and SAN EvaluScales, products are evaluated based on the following 4 criteria categories:

- Performance / Capacity
- Basic Functionality
- Advanced Capabilities
- Ability to Execute
-

The full NAS EvaluScale can be found [here](#).

The full SAN EvaluScale can be found [here](#).

The Futurum Group Opinion and Outlook for NetApp AFF

NetApp AFF is a unified SAN and NAS systems where block LUNs are mapped to files within the system and written to attached storage devices as blocks. As one of the first truly unified block and file storage systems, NetApp established the market for customers that want to use systems for both types of access. NetApp ONTAP systems, such as AFF, have an extensive set of advanced features including FabricPool tiering, SnapLock, and MetroCluster, and should meet the feature set needs of most organizations.

The AFF systems additionally provides flexibility as a block or file system, with additional S3 support for object storage, however object functionality may be limited compared to more mature object systems such as NetApp StorageGRID. The AFF line consists of several models ranging from entry level to high end enterprise systems, as well as designated C series systems that utilize QLC devices for capacity optimization, and ASA systems that are optimized for SAN only requirements. The wide range of models and support for NVMe devices and NVMeoF

should additionally help organizations meet most performance requirements. The AFF system is additionally a very scalable solution due to its scale-out capability, however it should be noted that the scale-out limitations are different for SAN and NAS uses.

Future outlook for NetApp and the AFF line looks strong. NetApp is a large and reputable vendor that has distinguished themselves as a leader in unified block and file storage. Its portfolio of ONTAP based storage offerings, including AFF, appear to be a core focus for the company and receive consistent updates and product enhancements. NetApp has expanded ONTAP to be the basis for their hybrid-multi-cloud strategy, with ONTAP offerings available in the public clouds and BlueXP used as a unified manager. NetApp AFF will remain the key on-premises offering as part of these hybrid environments.

Copyright 2023 The Futurum Group, LLC. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of The Futurum Group Inc. The information contained in this document is subject to change without notice. The Futurum Group assumes no responsibility for errors or omissions. The Futurum Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall The Futurum Group be liable for any indirect, special, consequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages.

ⁱ “File System Design for an NFS File Server Appliance”; Dave Hitz, James Lau, & Michael Malcolm ; NetApp; TR 3002