NetApp FAS3200 and V3200 Storage Devices

NetApp

Reference Code: IT017-003805
Publication Date: 05 November 2012
Author: Tim Stammers

SUMMARY

Catalyst

The FAS3200 and V3200 series are NetApp's mid-range line of disk arrays and storage controllers that were recently updated to increase performance and improve manageability. Like the rest of the FAS and V-Series family, they are unified devices that combine block-level (SAN) and file-level (NAS) storage in one device. The devices are sophisticated and are especially competitive in the area of storage efficiency, clustering and flash integration, maximizing performance, scaling, and utilization of physical resources.

Key findings

- Strengths include unified block and file storage architecture, multiple flash integration options, clustered scale-out to over 25PB, background data de-duplication and compression, and lightweight snapshots.
- Dual-controller architecture is less easy to scale than some rival multi-controller architectures, but FAS3200 devices can be upgraded non-disruptively to high-end FAS6200 disk arrays. NetApp does not disclose prices, but claims that FAS and V-Series 3200 devices are competitively priced.
- Competitive use of flash memory, with more options for deployment than other devices. FAS and V-Series 3200 systems can use flash in three ways, each with different characteristics to suit different applications.

Ovum view

IT departments should strongly consider the NetApp FAS3200 and V3200 series as a unified storage platform.

Value proposition

In 2002, NetApp became the first supplier to embrace the concept of unified storage, when it added block-level support to its storage devices, which had previously only allowed file-level access to data. It was not until 2011 that EMC became NetApp's first major competitor in unified storage, but the trend is now established, and other suppliers have begun to follow suit.

The volume of data that is handled as files, such as documents, images, and videos, has been growing very quickly for the last decade, and now accounts for the majority of business data. Creating shared storage for end-users is far simpler with NAS than block-level storage. NAS also simplifies the storage of virtual storage images. Applications such as Microsoft Exchange, SharePoint, and SQL Server only support SAN or block-level storage connected to servers using iSCSI or Fibre Channel. As a result, unified storage that supports those block-level protocols alongside file-level data access greatly simplifies operations for many customers.

Uniquely, for a device from a major supplier that is designed to store primary working copies of data, the FAS3200 series features block-level data de-duplication. This, in combination with data compression and a high-performing version of RAID 6 called RAID-DP, gives the FAS3200 series the potential to achieve high levels of efficiency in terms of disk usage. Finally, NetApp is maintaining a leading position in the use of flash to boost application performance, and lower purchase and operating costs. The FAS3200 series is a very competitive mid-range product.

SOLUTION ANALYSIS

Overview

The key parameters of the new FAS3200 series are:

- Maximum raw capacity of 2160TB across 720 drives in a single system
- 2.5-inch and 3.5-inch SAS performance disk drives
- 3.5-inch SATA capacity disk drives
- 3.5-inch SSD flash memory drives
- Up to 2TB PCIe flash memory extension to cache
- Third-party server-side flash devices (either PCIe Flash or SSD) integrated via NetApp software for cache coherency
- Aggregate internal pools of SSD flash and disk drive storage
- Active-active, dual redundant controllers, plus clustering up to 24 nodes for long term scaling
- Application-aware, reservation-less, and writable snapshots
- Synchronous data mirroring and distance-clustering, long-distance, asynchronous, periodic replication and replication-based data backup

- Up to 52 host ports, supporting Gigabit and 10GbE iSCSI and FCoE, and 2, 4, or 8G FC
- CIFS and NFS file-level access
- Non-volatile, high-availability memory to protect data writes, redundant power supplies, fans, battery backup, and call-home remote diagnostics
- 3U to 6U form factor.

The updated FAS3200 series comprises two models, which range in maximum capacity from 1TB to 2,160TB. The series is itself part of a wider family of NetApp FAS storage devices which comprises an entry-level FAS2200 series, the FAS3200 series, and a top-end FAS6200 series.

All of NetApp's FAS devices are based on the same hardware architecture and operating system, and are managed using the same software tools. The operating system is called Data ONTAP, and it includes a file system called Write Anywhere File Layout (WAFL). WAFL is one of the jewels in NetApp's technology crown. In the past NetApp referred to WAFL as a file system, but now the company calls it a "data virtualization layer."

WAFL includes a mechanism that boosts performance when handling data at both block and file level. The mechanism aggregates random writes so that they are completed sequentially. This accelerates write performance by eliminating the latency or "seek time" of disk-drive, read-write arm movements. It also improves disk-read performance, because data that has been written sequentially displays temporal and spatial locality, which again eliminates arm movements. A form of write journaling protects data against loss if an uncontrolled shutdown, such as loss of power, occurs before the aggregated writes have been written to disk.

In 2010, NetApp launched version 8.0 of Data ONTAP, which introduced the ability to cluster up to 12 HA pairs of FAS devices, comprising 24 controllers or nodes, in order to increase performance and capacity. Since then the company has progressively updated this clustering feature, which now allows hardware upgrades to be completed without taking a cluster offline or interrupting applications. This also allows non-disruptive load balancing to maximize performance for specific workloads.

The FAS3200 series was introduced in 2011, replacing the previous FAS3100 series. The most recent update was released in November 2012, and included the latest Intel multi-core processors, and an expanded memory footprint, which increased the throughput, capacity and maximum number of I/O ports of the devices.

The V3200 series of open storage controllers use the same hardware and Data ONTAP operating system as the FAS3200. However, the V-Series systems can front end third-party arrays from EMC, HP, Hitachi, and other suppliers, in addition to NetApp disk shelves. This enables customers to use older third-party arrays as though they were NetApp systems, extending the working life of those arrays.

Efficiency

Over several years, the relentless growth in the volume of data stored by businesses has spurred storage vendors to develop ways to reduce the amount of disk capacity needed to handle a given amount of data. As well as lowering customers' purchase or capital costs, these technologies also reduce operating costs, by lowering space, power, and cooling requirements. Collectively they are sometimes referred to as optimization or efficiency technologies. One of the most successful of these technologies has been thin provisioning. The FAS series has supported thin provisioning for several years, but the feature is now common in enterprise disk arrays, albeit with varying effectiveness.

Currently NetApp has strong claims to leadership in three other optimization areas, namely background data de-duplication, compression, and RAID efficiency. NetApp claims that collectively its efficiency features typically reduce disk requirements by 50% or more.

Block-level de-duplication and compression

NetApp added block-level data de-duplication to the FAS series several years ago, and is still the only major supplier offering this feature on disk arrays that were designed to store primary data, rather than being specialized backup targets. NetApp's block-level de-dupe applies to data accessed through both NAS and SAN. Some rival primary storage devices also de-duplicate data, but only at the file level, and with far less reduction in data volume.

In 2011, NetApp added data compression to the FAS series. It is not the only major vendor to offer this for primary working data, but the feature is still not common. Compression and de-duplication are complementary technologies. Not all data de-duplicates well, and the same is true for compression. The FAS3200 allows both optimization methods to be applied to data, either individually or in combination. An example of where de-duplication and compression provide the best combined savings is file shares. User file shares usually see around a 30% reduction in volume with de-duplication, a 50% reduction with compression, and a 65% reduction when both are used. In contrast, Oracle databases that write in 8KB blocks do not benefit from de-duplication, but NetApp claims that they commonly get 70% savings from compression. Independent of the FAS3200 series, NetApp provides a tool that examines data and predicts the savings that can be achieved by de-duplication and compression, separately and in combination. The data does not have to exist on a NetApp system to be analyzed.

Both compression and de-duplication run as low-priority background processes. De-duplication can be scheduled to run either on a time basis or trigger basis. Compression can be set to occur either before or after data is written to disk.

Resilient, fast, and space-efficient RAID

The FAS series involves a unique RAID strategy that contributes to the overall efficiency of disk space usage, while boosting data protection and maintaining performance. RAID data striping protects data against disk drive failures, and it can increase performance. It can also reduce performance, depending

on the RAID level or variant, but that loss is accepted as a trade-off for data protection. Generally, the greater the performance or protection offered by a RAID level, the more drive capacity it consumes. As a result, almost all enterprise disk arrays offer a choice of RAID levels, which are usually RAID 10, 5, and 6. This allows customers to choose the level that best meets the needs of an individual application, and the type of drive involved.

That choice is always a compromise between performance, data protection, and space consumption. In contrast NetApp recommends a single default RAID level for all FAS series applications, which offers high performance and data protection, as well as low space usage. This RAID level is called RAID-DP. NetApp has underlined its confidence in RAID-DP's suitability for all applications, by using it for all of its FAS performance benchmark tests.

RAID-DP is an implementation of RAID 6, but NetApp claims that it is significantly faster than other vendors' versions of RAID 6, and nearly as fast as RAID 5. This is because RAID-DP completes dual-parity, disk-write operations in parallel, unlike other RAID 6 implementations. In addition, RAID-DP benefits from the write acceleration of WAFL. Because RAID-DP involves dual-parity protection, it offers much higher protection than RAID 5. It also uses slightly more disk space than RAID 5. NetApp has partly addressed this issue by allowing RAID-DP to work across large groups of disk drives, improving its space efficiency.

Other than RAID-DP, the FAS series only supports RAID 4. Although RAID 4 uses less disk space and is 2% faster than RAID-DP with write workloads, it does not offer the same level of data protection. NetApp only recommends RAID 4 for a minority of "scratch space" and other applications that require less protection.

RAID 10 is generally faster than other RAID levels. As a result, it is widely offered on rival disk arrays, for use with performance-sensitive applications. NetApp acknowledges that RAID 10 might be a useful option for the FAS series. However, the company states that FAS performance using RAID-DP already exceeds customer requirements. Also, RAID 10 uses considerably more disk space than RAID-DP, but provides less data protection. As a result, NetApp currently sees no need to offer RAID 10.

"Zero impact" snapshots

A standard feature of mid-range and high-end disk arrays is the ability to create space-efficient snapshot copies of data, which are integrated with a range of commonly used applications so that they can be used for backup. The FAS series can do this, and it can store many more snapshots than its rivals. This quality is especially beneficial when exploiting the virtues of snapshots as a backup mechanism. Unlike traditional backups, snapshots can be created almost instantaneously, and therefore frequently, during working hours. When data has to be retrieved from a backup, recovery from a snapshot is far quicker than from a conventional backup.

Because snapshots only record changes to data blocks, they consume a fraction of the disk space of the source data from which they were created. The more snapshots that are kept on an array, the

longer and more frequent the history of snapshots available for roll-back recovery to a previous point in time. However, for some disk arrays, storing large numbers of snapshots heavily reduces overall performance. This is not the case for the NetApp FAS series. FAS devices can maintain up to 255 incremental snapshots of every volume they store. At 1,000 volumes per HA pair of systems, that means that the FAS3200 supports up to 255,000 snapshots per HA pair. As those snapshots apply a negligible load, NetApp recommends that its customers fully exploit this capacity. Other vendors may also quote large figures for maximum numbers of snapshots, but only at the penalty of heavily reduced performance.

This is because snapshots record changes to data blocks using pointer systems, which continually consume I/O and other resources during the life of the snapshots. The FAS series does this using an efficient algorithm called redirect-on-write. This was implemented by NetApp several years ago, but is only now beginning to appear in rival products. Even compared to the few rivals that have made the move to redirect-on-write, NetApp claims that WAFL allows its snapshots to apply less load.

Flash Integration

Over the last three to four years, flash memory technology has become a standard and widely used option in mainstream storage systems, because of its major performance and cost benefits. Adding a relatively small amount of flash to a disk array not only boosts its performance, but also reduces its upfront purchase costs, by allowing low-cost high capacity disk drives to be used in place of expensive, fast disk drives. Flash usage also reduces power, cooling and space requirements.

This first wave of flash adoption is now being followed by the development of new and more sophisticated ways of integrating flash into storage systems. This is a competitive field, in which NetApp is maintaining a strong position. Unlike many of its rivals, NetApp now offers three ways to use flash, each with characteristics that suit different applications, and all under the label Virtual Storage Tiering.

Flash Cache

This option allows FAS storage systems to be fitted with PCIe Flash cards that are connected to FAS controllers. Data ONTAP automatically loads the flash memory with a cache copy of whatever is currently the “hottest” or most frequently accessed data. Flash Cache is architecturally a little different from other vendors’ so-called automatic tiering systems, which move hot data wholesale from disk into flash, where it is both read from and over-written. In contrast, Flash Cache is used only for reading data, and so accelerates random read operations but not write operations. NetApp says that this is the more efficient approach because it avoids wholesale data movements, and because of the WAFL virtualization layer within all FAS and V-series devices. WAFL accelerates write operations, and is unique to NetApp.

The effects on performance and cost of the two competing architectures, caching or wholesale data movements, vary according to applications and system configurations. NetApp claims that Flash Cache can boost performance for applications such as Microsoft Exchange and desktop virtualization by 75%,

while reducing upfront purchase costs by 50%. These numbers are very similar to the percentages quoted by other suppliers for the use of flash with conventional tiering functions that move data wholesale.

Flash Accel

NetApp's Flash Accel feature implements so-called server-side caching of data. This architecture is a very recent industry development, and NetApp is one of the first major storage suppliers to offer it as an integral part of its products. Flash Accel performs a very similar function to Flash Cache, but instead of maintaining a cache of data within a FAS storage system, it maintains the cache within the servers that are accessing that data. This increases performance by eliminating the latency or time taken for the server to communicate with the storage system when reading hot data. It also reduces the load on the storage system itself, boosting its performance when handling other data. Flash Accel suits applications with very high performance requirements, and NetApp says it can reduce latency by 90%, and boost overall system throughput by 80%.

Flash Accel works with a wide range of third-party PCIe or SSD flash drives installed in customers' servers. The NetApp product itself is software that runs on those servers, and identifies the hottest or best data to copy into the server-side cache. Additionally, Flash Accel promises greater overall value than the use of stand-alone server flash drives as direct attached storage, since it allows flash to be presented as a shared, network accessible resource. Flash Accel currently supports VMware's vSphere virtualization platform, including VMware's vMotion live migration of virtual servers. Flash Accel is also integrated with Data ONTAP, which minimizes the amount of cache flushing required to maintain data integrity, when data held within a FAS system is changed by snapshot operations. When this happens, other systems must flush entire caches, reducing their performance during the subsequent and sometimes lengthy process of "re-warming" the cache, which involves re-identifying and then re-loading the hottest data back into cache.

Flash Pool

The third VST implementation is Flash Pool, which is a Data ONTAP aggregate that consists of a combination of SATA or SAS-attached disk and flash drives. Within these aggregate allocations of storage capacity, Data ONTAP automatically moves the hottest data to flash, in small and efficient 4KB blocks. Like Flash Cache and Flash Accel, Flash Pool accelerates random read operations. Unlike those two other flash mechanisms, Flash Pools use RAID-protected flash to handle writes as well as reads, which boosts the performance of random–write intensive applications such as OLTP. Because the flash is part of the storage system, Flash Pools provide fast recovery from unplanned controller outages, caused by events such as power failures.

Scalability

As data continues to grow in size, the ability to expand the capacity of a storage device without lowering its performance has become crucial for customers. This is scalability, and in one important aspect the

FAS series is arguably one of the most scalable storage devices on the market, because the same architecture, operating system, and management tools power a range of FAS devices, spanning entry-level to high-end enterprise machines.

The architecture involves a traditional dual active-active controller setup. Unlike some rival devices that are powered by multiple controllers, this does not allow FAS data processing power, and hence capacity for a given performance, to be increased by simply adding controllers. Instead FAS series devices are upgraded by replacing the controllers, in order to step up from one FAS series to the next. Doing this involves shutting down the controllers one at a time, with resulting effects on performance during the upgrade, and risk of problems during the conversion. However, the upgrade only requires the controllers to be replaced, as the disk shelves can remain in place. In other words, it is a partial, in-place upgrade, not a forklift upgrade.

Alternatively, multiple FAS units can be tied together into a scale-out cluster of storage systems, which share access to the same data. NetApp argues that this is more flexible than rival systems that can be scaled up simply by adding controllers to one device. In rival multi-controller devices, controllers are embedded in disk shelves. This means that capacity can only be increased by also increasing the number of controllers and network ports, whether or not extra controller processing power and input/output (I/O) are needed. For a cluster of separate FAS systems, disk shelves can be added for capacity, independent of the controllers and without consuming extra I/O ports. As a result, clusters of FAS systems can be scaled up to meet application requirements by capacity, CPU performance, or I/O ports, independently in each parameter. Data stored in a cluster can be migrated between nodes during system upgrades, or load-balancing exercise, non-disruptively without affecting applications.

Clusters of up to 12 HA pairs of FAS systems can be created, comprising 24 nodes or controllers. For the 3200 series systems, this can produce a cluster with a total capacity of over 25 petabytes. In addition, a single volume within the cluster can be as large as 20PB.

Performance

NetApp cites benchmark test results that indicate impressive performance handling data at both the file and block level. The block-level results come from tests designed and supervised by the Storage Performance Council (SPC), which are widely used benchmarks. The SPC tests include two types of test for disk arrays – one that is intended to simulate random I/O loads, and one that is intended to test sequential throughput. The file (NAS) oriented benchmarks are designed and supervised by the Standard Performance Evaluation Corporation (SPEC).

All of NetApp's SPC block-level tests were completed to the benchmark for random I/O performance, rather than the SPC's alternative benchmark for sequentially biased I/O performance. This decision by NetApp is a good indication that even though the FAS series began life as a file-level, NAS device, it now also performs well when handling database and other block-level applications that generate the most challenging random I/O. NetApp points to WAFL as the reason for this. One caveat for customers is that benchmark tests can only be approximations of real-world performance. Another is that vendors

sometimes complete these tests using unrealistic racing configurations of their products. For all SPC results, customers should pay attention to the price of the configuration per achieved SPC IOPS rating. This is calculated and listed as part of the official SPC results.

Management tools

Similarly to other enterprise disk arrays, the FAS series is supported by a wide range of optional software tools providing snapshot vaulting, replication, and other functions. The quality of these tools has been a major reason for NetApp's success over the last two decades, and the company has been continually broadening their functional scope.

Data ONTAP Essentials management tools are standard with all FAS devices. OnCommand is the family name for the management software products, which include a web-based element manager, OnCommand System Manager, that NetApp says can be used with no specialist knowledge. OnCommand Unified Manager combines automated monitoring with data provisioning and protection, and is recommended for organizations running more than a few FAS units.

OnCommand also includes cross-vendor analytic tools that were originally developed by storage specialists Onaro and Akorri, which NetApp bought in 2008 and 2011 respectively. OnCommand Insight (formerly Onaro) identifies performance bottlenecks, monitors application response times and disk utilization, and helps with capacity planning. OnCommand Insight also reports on application-level usage of heterogeneous resources, from virtual infrastructure through to storage. OnCommand Balance (former Akorri) provides agent-less performance management for shared infrastructures. It correlates and analyzes data from VMs, physical servers, and storage for troubleshooting, optimization, and prediction capabilities.

Data ONTAP also includes:

- MetroCluster, a distance clustering solution that extends high availability beyond the data center and protects against both planned and unplanned downtime.
- SnapMirror, a data replication mechanism that uses bandwidth-efficient, thin replication to create copies of snapshots on other FAS or VSeries devices.
- SnapVault, a disk-backup tool that moves snapshots to another FAS system where they can be stored for weeks, months or years, while maintaining their ability to provide rapid data recovery.
- SnapLock, which creates immutable WORM data volumes to meet legal hold and regulatory requirements.

NetApp snapshots can also be managed by traditional backup tools, namely Symantec NetBackup, CommVault Simpana, and Syncsort data protection software. This allows snapshots to be integrated with traditional backup processes.

Go-to-market strategy

Since it began in 1992, as a pioneer of what was then the new concept of NAS, NetApp has grown quickly. For the five years to 2011, NetApp claims 20% CAGR in revenue, and states that the only other IT companies that have exceeded $4bn in revenue while also showing five years of 20% organic growth are Amazon, Apple, eBay, Google, and Huawei.

NetApp's growth has occurred partly because the company was in the right place at the right time. Having recognized the potential value of NAS at the beginning of the 1990s, NetApp has since enjoyed growing demand for storage to handle the ballooning volumes of file-level data stored by businesses. Much of NetApp's success has also been the result of its product. Other suppliers, including OEMs such as HP and IBM, eventually recognized the value of NAS, and attempted to compete with NetApp. However, NetApp still only has one major competitor, EMC, in NAS, and more recently in unified storage. In 2005, IBM decided that the best way to expand its NAS offerings was to resell re-branded NetApp products, rather than develop its own. That deal has worked well for IBM, and is still in place.

NetApp is now consistently among the five largest suppliers of disk arrays according to market share, and has broadened its product range to include object-oriented storage systems and block-level only systems. However the company faces the long-term threat of being made irrelevant by supplier consolidation, which is being driven by the possibility that technology convergence will force customers to seek single suppliers of entire infrastructures. One response to this convergence has been NetApp's alliance with Cisco and VMware, in which the trio sell ready-configured systems called FlexPod. These comprise Cisco's networking gear and servers, NetApp's storage, and VMware's virtualization software.

NetApp is confident that it will survive the convergence trend by continuing to sell best-of-breed products. During the last decade IBM, Cisco and, Oracle have been periodically rumored to be negotiating to buy NetApp. NetApp says it is not interested in selling, and points to its giant market valuation, $16bn currently, as a reason why, in reality, acquisition offers are unlikely. Even if NetApp was bought, its product lines would continue to enjoy a strong future.

Deployment

NetApp claims a pilot deployment of a FAS3200 device typically takes two hours to one day to achieve, and can be completed by one person. A 30-user departmental deployment would take between half a day and three days, using one to three full-time staff. A 500-user enterprise deployment would take up to five days, using two to six staff.

A 30-user departmental deployment requires basic knowledge of storage and networking concepts, including iSCSI or CIFS. Knowledge of NFS and Fibre Channel may also be required. NetApp says that knowledge of application integration features, for example between the FAS series and VMware, Oracle, or Microsoft Exchange, would be helpful but is not essential. Enterprise-wide deployments will

require further knowledge of NetApp storage management. Advanced knowledge of technologies such as Fibre Channel, FCoE, Gigabit, and 10GBit Ethernet may also be required.

Training for simple deployments can be completed using NetApp web-based training tools. More complex deployments may also require instructor-led training. For very large and complex deployments, customers typically send two or more administrators for one week of NetApp classes, plus extensive online training in specific storage management topics.

NetApp sells products directly itself, and indirectly through resellers and integrators. NetApp-authorized service partners span from global organizations such as Accenture, IBM and Wipro, to value-added resellers and industry-specific or vertical solution integrators. Services range from basic “rack and stack” to full deployment and integration with applications and server virtualization systems, as well as assessment services for data migration and storage consolidation projects.

Software is predominantly licensed per system. A basic level of support is included free with all FAS3200 and V3200 devices, and comprises a three-year hardware warranty with next- business-day delivery of replacement parts, phone home “auto-support,” monitoring and diagnostic tools, and access to a support website. However, NetApp says that nearly all customers choose to buy one of two more advanced levels of support, the most advanced of which provides replacement parts within two hours, and “targeted” response times of 30 minutes.

DATA SHEET

Key facts

Data sheet

Product name	FAS2000 and V3200 series	Product classification	Storage device
Version number	NA	Release date	2011
Industries covered	NA	Geographies covered	Global
Relevant company sizes	Usually around 1,000 or more employees	Platforms supported	Windows - All versions for CIFS. All recent server versions for iSCSI and Fibre Channel. Linux - All versions for NFS. More recent Red Hat and SUSE versions for iSCSI and Fibre Channel. Solaris - All versions for NFS. More recent versions for Fibre Channel. AIX - Last five years' versions for NFS. More recent versions for iSCSI and Fibre Channel. HP/UX – Last five years' versions for NFS. More recent versions for Fibre Channel. z/OS – Support for NFS Mac OS - Most versions of OSX (10.x) for CIFS and NFS. Limited support for Fibre Channel IBM i-series (OS/400) - NFS-only support.
Languages supported	NA	Licensing options	NA
Deployment options	NA	Route(s) to market	Direct and channel
URL	www.netapp.com	Company headquarters	Sunnyvale, California
European headquarters	Amsterdam	North America headquarters	Sunnyvale, California
Asia-Pacific headquarters	NA		

Source: Ovum

APPENDIX

Further reading

Ovum Technology Audit, EMC VNX Storage Device, May 2011

Methodology

- Briefings with vendors
- Comparison of multiple vendors' technical literature.

Author

Tim Stammers, Senior Analyst, Infrastructure

tim.stammers@ovum.com

Ovum Consulting

We hope that this analysis will help you make informed and imaginative business decisions. If you have further requirements, Ovum's consulting team may be able to help you. For more information about Ovum's consulting capabilities, please contact us directly at consulting@ovum.com.

Disclaimer