Creation: Jan. 23, 2013
Modified: March 5, 2013
Jeremy Li

Evaluation of NetApp Storage V-Series Product:

EXECUTIVE SUMMARY

Although NetApp sells its hardware based storages, it claims NetApp is a software company that has about 40 different applications and licenses (Click http://lacaaea.com/vendors/sw-lic.pdf) for details) any enterprise needs to purchase in order to utilize its efficient thin snapshot and replication technology effectively. For example, an enterprise needs to purchase (1) a license to take a snapshot; (2) a license to replicate snapshots; (3) a license to be able to recover a snapshot at the remote site. When a final price includes the software licenses, an enterprise always finds that the price jumps. Customers know they can get the same storage capacity and similar software features for less money elsewhere.

NetApp V-Series solution might be a good combination solution for enterprise(s) because it can be complement to the EMC's inefficient software-based snapshots due to its Copy on Write (COW) technology.

The famous NetApp Thin Replication consists of **SnapMirror** (transfers only Changed blocks, Deduplicated data and Compressed data), **SnapVault** (With the combination of SnapMirror and SnapVault, only incremental block changes are transferred as long as the baseline copy is made, as illustrated in the screenshot below), **SnapProtect** (Disaster Recovery/Business Continuance Applications) and **FlexClone (**Per NetApp, no test & develop environment should be without it). For example, in a real world environment, FlexClone® and Snapshot reduced development & test environments with capacity from 160TB to 18TB, and clone creation takes less than 2 minutes, restore takes less than 2 minutes, as illustrated in the picture below.

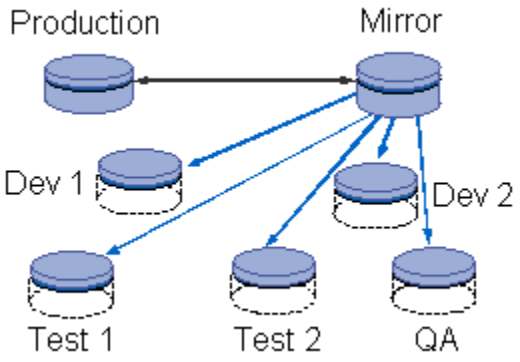

SnapMirror® Replication only transfers 17GB data instead of 200GB data between two long distance sites for disaster recovery purpose. As a result, a bandwidth upgrade between sites is avoided. Its snapshot typically can reduce the restore time from days to hours with up to 80% tapes eliminated.

In VMware primary storage:

- 86% deduplication rate
- Combination of deduplication and Snapshots freed 4TB in a 6TB storage array

Generally speaking, an enterprise needs to purchase a clustered pair with NetApp Disk Shelves. 30% of all V-Series with some NetApp disks will be purchased at the initial acquisition, while 66% have some NetApp disks over time. Flash Cache: 57% of V32XX and V62XX have Flash Cache.

However, EMC does have a high quality snapshots and replication products via its hardware based solutions at a high price. It provides continuous local and remote protection with application consistent bookmarks capability.
EMC has two hardware-based replication products:

1. Symmetrix Remote Data Facility, better known as SRDF, can replicate date from one Symmetrix storage array to another through a Storage Area Network (SAN) or IP network (EMC Products only).
2. RecoverPoint, as illustrated in the picture on the right  , can replicate data from EMC storage to any storage including a few third party storages. It can enable any point-in-time recovery.

Selecting a hardware-based replication or a software-based replication really depends on an environment. EMC is still #1 storage vendor with highest market share due to its reputation and superior technical support. IBM selected EMC (VMAX family) as OEM for its System z and IBM i, while IBM selected NetApp to supply its mid-range NAS products. Therefore, choosing EMC will not let you go wrong, as long as a capital investment is not an issue. On the other hand, an enterprise often chooses a software-based thin replication technology because a hardware-based replication aren't economically feasible in order to deliver a similar result for disaster recovery (DR) and Business Continuity Plan (BCP) due to avoidance of additional layer of complexity.

NetApp uses its highly efficient snapshots from its OTAP OS to achieve the DR and BCP objectives for many enterprises. It gained 60% market share in Federal Government Agencies as of 2012. Based on reading some of the documents (e.g., JCVI Supports Groundbreaking Genomics Research with NetApp Storage and ESG Lab Validation Report) supplied by NetApp and listed under the appendix A, a question and answer session was conducted from 11:00 a.m. to 1:00 p.m. on January 23rd with a NetApp engineer.

Q1: Does the V-Series Solution Pass the Validation Test by EMC?

A: NetApp purchased all EMC storages, devices and went through the full validation in house. The compatibility is not an issue between two products as long as firmware versions are met.

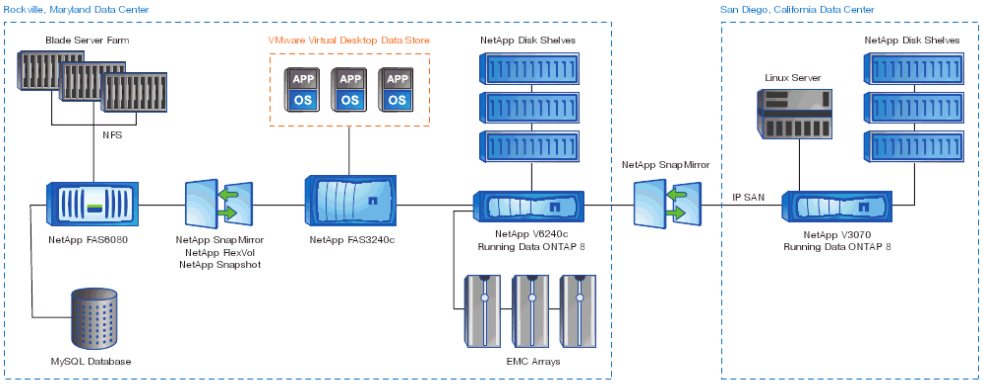

Figure 1) JCVI storage infrastructure.

Click http://www.netapp.com/us/system/pdf-reader.aspx?m=jcvi-ss.pdf&cc=us for details.

Q2: Advantages and Disadvantages of NAS

The following presentation slide was displayed on the projector screen in Microsoft training class titled MS-20410 "Installing and Configuring Windows Server 2012" as well as in the courseware.

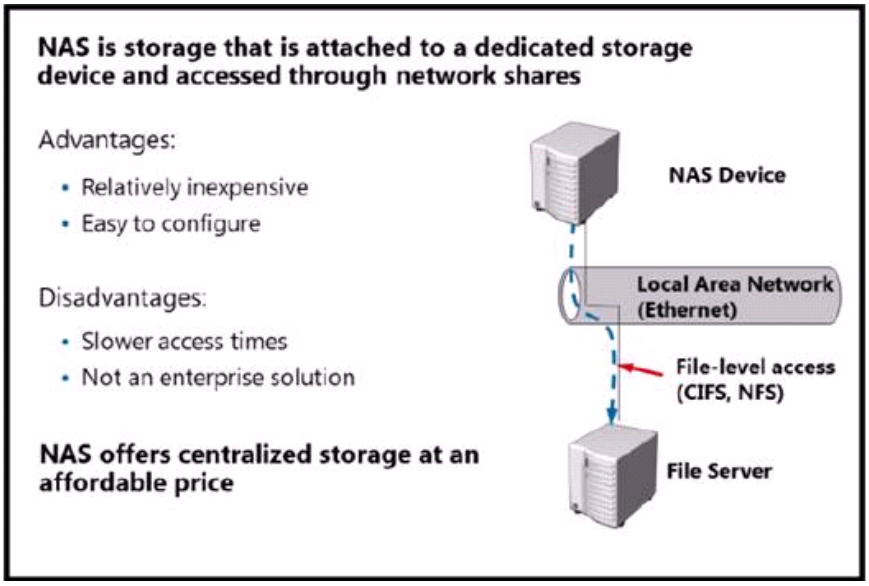

A: NAS is simple Enterprise solution with faster access times than Windows or Unix File server systems. Why is Microsoft moving key applications back to NAS/CIFS from block based protocols? You need to ask yourself a question: Why are NetApp Enterprise customers moving their VM's to NAS/NFS from FC and iSCSI? Pretty simple, it is simple and fast.

Author's comments:

The 2012 storage report from Goldman Sachs listed below shows that NetApp ranks as a top share gainer: (http://executivecouncil.com/reports/2012_July_GS_CIO_Report.pdf)

Storage: NetApp ranks as a top share gainer, while EMC's streak continues

Exhibit 14: Which storage providers are gaining or losing share of your IT spending dollars?

Gaining	Losing
1. NetApp	1. HP
2. EMC	2. IBM
3. Dell	

Source: Goldman Sachs IT Spending Survey.

However, EMC is still #1 storage vendor with highest market share due to its reputation and superior technical support. . IBM selected EMC (VMAX family) as OEM for its System z and IBM i. Choosing EMC will not let you go wrong, as long as the capital investment is not a question.

Q3: Does NetApp Support Exchange Server 2010 in a Virtual Environment?

Here are quotes from other Storage Pros:
NetApp was not created for open systems. They were designed for NAS file.
You can research it and find that they should not even be used on Exchange in a virtual environment. This is an excerpt from one of their recent press releases.

- We continue to believe that NetApp's one-system-fits-all message is too narrow, representing a stark contrast to EMC (EMC $23.40; Outperform), Dell (DELL $9.58), IBM (IBM $185.51), HP (HPQ $13.14), and others who offer multiple platforms for different use cases. It is also evident that EMC and emerging players such as Nimble Storage and Nexenta continue to encroach on NetApp's midrange footprint, a segment that represents the company's historical stronghold. In our view, NetApp has struggled to capitalize on some of the fastest-growing segments of the storage market, including flash (hybrid and all-flash arrays, and server-side PCIe) and data protection.

- An excerpt below from Microsoft TechNet (Covers both Exchange 2007 and Exchange 2010).

"All storage that is used by an Exchange guest machine for the storage of Exchange data must be block-level storage. Exchange 2010 does not support using network attached storage (NAS) volumes. NAS storage that is presented to the guest as block-level storage by using the hypervisor is not supported. Pass-through volumes must be presented as block-level storage to the hardware virtualization software. This is because Exchange 2010 does not support using network attached storage (NAS) volumes. The following virtual disk requirements apply to volumes that are used to store Exchange data."

The full article can be obtained from the link below, and the quoted text motioned above can be found while expanding the "Support Policy and Recommendations for Exchange Server 2007" section.

http://technet.microsoft.com/en-us/library/cc794548%28EXCHG.80%29.aspx

> Follow up to the fantasy football yahoo outage.
> http://m.bizjournals.com/sanjose/news/2012/11/15/mayers-recovery-after-fumbling.html?r=full

A: NetApp has thousands of customers run Exchange 2003/2007/2010 on its storage without any issue.

NetApp suggests running an Exchange application via NFS in one of three configurations Microsoft will support.

Below are three configurations recommended by NetApp:

1. NFS > VMware > Windows Host > iSCSI > Database

http://social.technet.microsoft.com/Forums/en-US/exchange2010/thread/1b3e88d3-b4e9-455e-8fef-bbad929d3ee2

2. FC > VMware > Windows Host > iSCSI > Database

Note: It may not be worth it for using an expensive FC investment, then, switching it back to iSCSI in order to get Microsoft support.

3. Fiber Channel > VMware > Windows Host > RDM > Database

Note: RDM stands for Raw Disk Mapping, a hard disk's mapping for use by a virtual disk image. Using RDM via Site Recovery Manager will be a nightmare. Therefore, most customers do not choose to use RDM.

Please note Microsoft will not support the following configuration on NetApp box:

4. Fiber Channel > VMware > Windows Host > NFS > Database

Q4: What is the Price Difference between EMC VNX/VMAX and NetApp E Series?

Whenever possible, please give me a general price difference between two vendors!
See Q16 for your reference.

A: NetApp does not want to discuss this issue. Per NetApp Account Executive, some enterprises will not use it, even if the NetApp storage is free of charge. It is a political issue.

Q5: Does NetApp Have a Forklift Upgrade Issue?

A: No. NetApp can change two controller headers to upgrade its existing storage in order to preserve customers' investment. For example, after 5 years' investment on NetApp storage, an enterprise can either upgrade or replace the old headers with more powerful headers to continue to use its existing storage or keep the old disks and purchasing newer controllers. In a common scenario, a storage array usually consists of 80% disk drives and 20% controllers. Therefore, simply replacing two controllers will preserve 80% of the original investment, instead of going through a forklift upgrade – buying a new storage enclosure or a new model.

Dell Compellent SC8000 or Model 40 is a champion for avoiding a forklift upgrade. However, when its two controllers reach a full capacity (100%), it will go through a forklift upgrade per NetApp (must buy a second array).

That's why NetApp owns 60% Market Share in Federal Governments, and is the second largest storage vendor behind EMC.

Q6: May You Tell Me More Successful Stories from NetApp's VIP Clients by Using the V-Series and Co-existing with EMC VNX/VMAX and HP 3PAR?

A: California DWP has been utilizing its highly efficient snapshots to replicate data to multiple sites without relying on a secondary storage backup (e.g., EMC data domain, a disk-to-disk backup appliance)

Special note: When the V-Series is deployed, NetApp guarantees its storage saving at 50% level via its Dedupe and thin provisioning technique. However, all existing data muse be moved from EMC storage to NetApp storage in order to perform the Dedupe task.

Q7: What's the File Size Threshold from NetApp DATA ONTAP 8.1 in order to Convert Any Non-usable Disk Space (Whitespace with 00000000) into a Reusable Disk Space or Volume?

For example, HP 3PAR claims that it can perform much efficient reclamation granularity with 16KB compared to 768KB with EMC VMAX, as shown below:

Note: NetApp Data ONTAP 8.1 adds the clustering capability. That immediately increases its storage capacity.

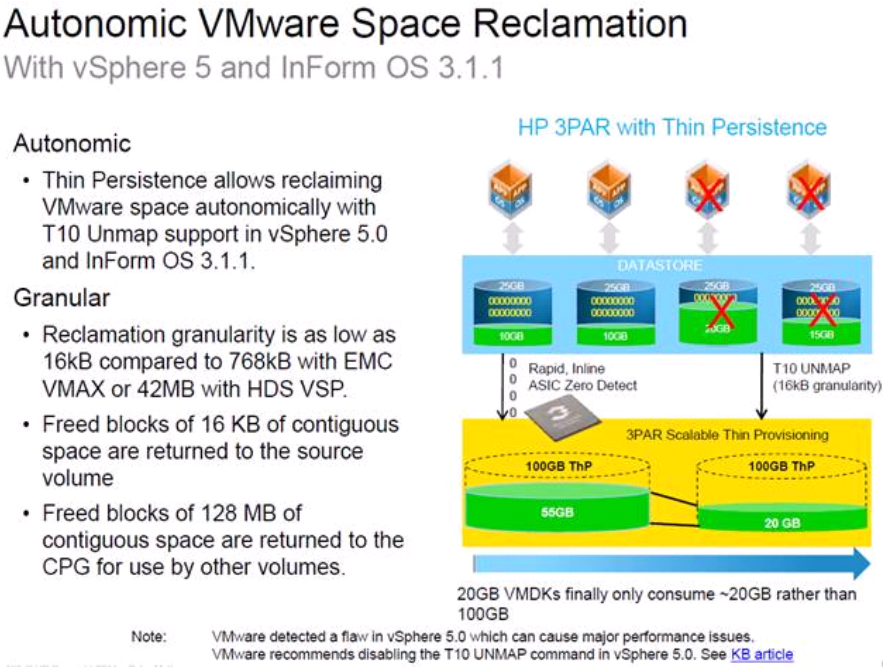

As a result, below shows that fewer HDDs will be purchased via HP 3PAR thin technology due to capability of eliminating extra zero (00000000) above 16KB size, instead of EMC VMAX 768KB. In other words, any file size smaller than 768KB, EMC VMAX can not reclaim or convert those wasted disk space into a reusable disk space. Below are the less physically installed disks after HP 3PAR Array's ASIC feature turned on.

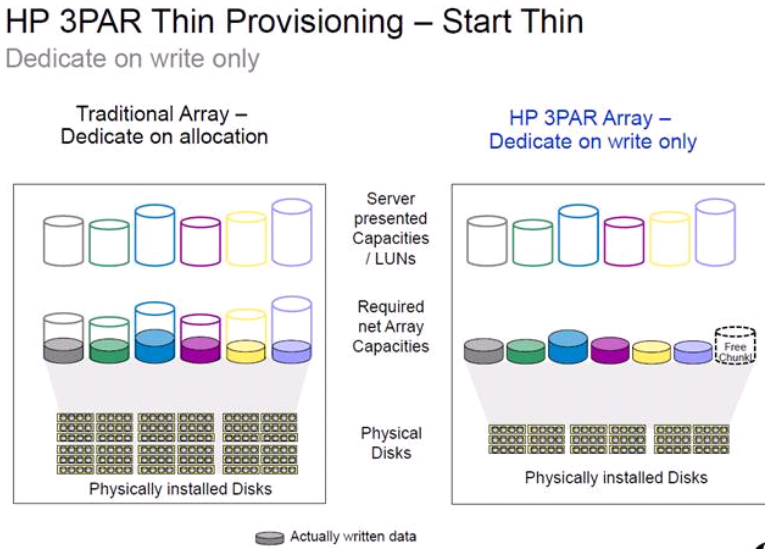

In summary, HP 3PAR's customers can get more free disk space up to 75% via its ASIC, as illustrated in a screenshot below:

HP 3PAR Thin Conversion – Get Thin

Thin your online SAN storage up to 75%

A practical and effective solution to eliminate costs associated with:

- Storage arrays and capacity
- Software licensing and support
- Power, cooling, and floor space

Unique 3PAR ASIC with built-in zero detection delivers:

- Simplicity and speed – eliminate the time & complexity of getting thin
- Choice - open and heterogeneous migrations for any-to-3PAR migrations
- Preserved service levels – high performance during migrations

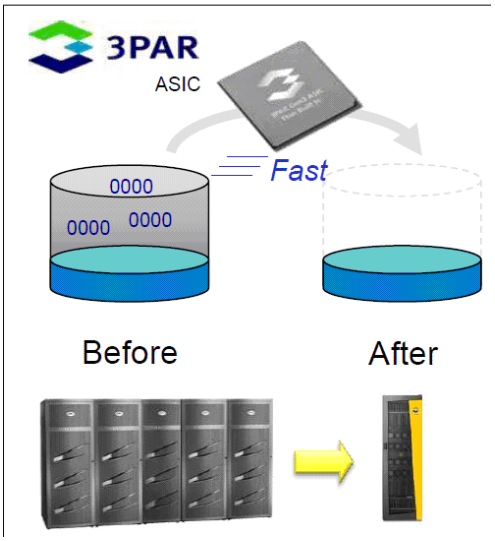

A: Per NetApp, it is useless, unless a FULL FORMAT is performed during a new volume creation period. Microsoft Windows OS does not write zeros to any storage or file systems. If you delete data, no zero will be created.

Note: Although author asked a HP Partner for comments and do not receive any feedback, he believes there are some values on this area. For example, it might be useful in Linux or Oracle environment.

If a file size is at 4KB, does it occupy a disk space at 768KB in VMAX storage?

Q8: May You Tell Me the Performance Penalty When the Dedupe is Enabled (A-SIS) in the NetApp Primary Storage?

Although Data reduction traditionally confined to backup systems (e.g., EMC Data Domain), it is becoming ubiquitou in the storage ecosystem and now common in Tier 1 storage arrays, storage virtualization software, cloud storage gateways, even SSD controllers. For example, how long will the deduplication be completed for 1000 TB data?

a) Most NetApp customers run the Dedupe on weekend (once a week) due to performance issue.
b) While the Dedupe is running in a background, can a daily backup be run?
c) Can the Dedupe be turned on or off via scheduling task? **Note:** I was told that once the Dedupe is turned on, you can not turn it off.

A: The maximum size of storage for the Dedupe is 100TB per volume. The Dedupe can be turned on or off and can be configured via scheduling task. Usually, it takes 8 hours to complete the Dedupe process for 100TB data.

Q9: Since NetApp Replication is Dedupe-aware to further Improve its Replication Efficiency, May You Tell Me about the EMC Replication Technology, including PROS and CONS between Two Vendors' Replication Technologies?
For example, most EMC customers rely on the following techniques:

1) VMware's SRM (Site Recovery Manager) to facilitate a failover
2) The EMC's SRDF (Symmetrix Remote Data Facility) to replicate data between datacenters

A: Enterprise must purchase two or four (e.g., SRDF) similar or identical devices in order to perform its replication, while NetApp can replicate to any non-identical or different array(s) to achieve its flexibility.

Q10: Can EMC VMAX and VNX Run Dedupe on a Primary Storage?

Note: EMC Data Domain (DD), a secondary storage, has a great feature of Dedupe. EMC won the biding war over the NetApp a few years ago to acquire the DD.

A: NetApp thinks the secondary disk-to-disk storage might not be necessary if enterprise(s) can utilize its highly efficient snapshots to replicate data between two sites or multiple sites.

Q11: What's the Overhead on NetApp WAFL (e.g., 30%) and RAID, Aggregates, FlexVols?

Since HP 3PAR and Dell SC8000 use a technique called "Disk Striping" with a block size of 1GB and 2MB, respectively, across all available disks (HDDs) within a system, while the WAFL relies on the fixed RAID group that contains 14/16 HDDs in most scenarios via a method called "Aggregates". This will restrict the performance IOPs dramatically. That's the reason NetApp and EMC customers must rely on CACHE (e.g., flash module) to address the disk bottle neck.

Note: Many other storage vendors told me that Aggregates method is 20 years old technology.
Reference: https://communities.netapp.com/thread/10591

With RAID DP you create multiple RAID groups of 14 disks, an aggregates is created on top of these groups and then FlexVols are created within this aggregate?

If I wanted to build an array with 100TB usable capacity (an example) would I create many RAID DP groups (max disk 14/16?) with many aggregates (max size 16TB?) then create lots of FlexVols within these aggregates then give hosts access to these FlexVols?

What is the purpose of the root aggregate and how and when should these be created? Or # of disks?

Do you need a hot spare per tray or can you have global?

Also, can anyone point me in the direction of good documentation, particularly with respect to RAID, aggregates, FlexVols and how they are put together with large arrays?

Many thanks for any help.

Chris Kranz Aug 12, 2010 1:37 PM (in response to nsitps1976)
I'll answer a few of your questions...
Yes you have disks, then RAID groups (14 for SATA, 16 for FC by default, this can be tweaked), then aggregates, then volumes, then data (or LUNs, then data).

If you want 100TB usable then yes, you do need to carve it up into separate FlexVols in order to provide it all. ONTAP 8 introduces 64bit aggregates, so the 16TB cap is raised according to the hardware you use.

The root aggregate is simply whatever aggregate contains vol0, this is the system volume that contains the main OS. If you have a system as large as you are saying, you would separate out the root aggregate to have its own disks. If you have a smaller system where disk is at a premium (2000 series), it can be within a normal data aggregate.

Remember that dedupe (A-SIS), clones and thin provisioning can allow you to allocate and address more storage than you technically have.

Your hot spares are global per controller, per disk type. So in a cluster you need a minimum of 2 hot spares per node. If you have SATA and FC disks, you'll need 2 of each. There is a table that shows how many spares to keep according to the number of disks you have in total, but I can't recall exactly where (hopefully another member will be able to help you out).

Have a look around http://www.netapp.com/us/library/ and look at the TR's. There are a couple of best practice guides. If you have a NOW account (now.netapp.com), search the production documentation libraries, these will show you the recommended maximums and minimums of a configuration.

A: NetApp can support 24 SAS HDDs and 20 SATA HDDs per RAID group. HP and Dell do the same trick via RAID Group (14+2). Click on the following link for details:
http://h18006.www1.hp.com/storage/pdfs/4AA3-3516ENW.pdf

Below are the excerpts from the above link:

RAID data protection

The HP 3PAR Storage System is capable of RAID 1+0 (mirrored then striped), RAID 5+0 (RAID 5 distributed parity, striped in an X+1 configuration where X can be between 2 and 8), or RAID MP (multiple distributed parity, currently striped with either a 6+2 or 14+2 configuration). The RAID 5+0 and RAID MP algorithms allow HP 3PAR to create parity sets on different drives in different drive cages with separate power domains for maximum integrity protection.

Simplifying the IT infrastructure requires that next-generation storage architectures provide consolidation, bi-directional scalability, and mixed workload support. The HP 3PAR Storage System addresses all of these requirements and provides multi-tenancy and autonomic management capabilities along with carrier-class availability that includes full software and hardware fault tolerance (See page 22).

Q12: What's the Rebuilding Time w/ 2TB Disk due to Dedicated Parity Disk (Bottleneck)?

For example, a Gridstore (http://www.gridstore.com/) 2TB appliance will take about **4.4** hours to rebuild a failed 2TB SATA disk if a 80% full capacity is reached via one 1Gbps NIC interface (1.6TB / 100MB per second = 4.4 hours). Click on the link below for more info:

https://subscriber.emediausa.com/FM/GetFile.aspx?id=34936.1.2818169.WXXLXCKH
NetApp and EMC use traditional arrays method to address the data recovery. Typically, 14 HDDs will be tied to one spare disk drive, while other vendors' rebuilding time will be greatly reduced due to Spare chunklets - distributed sparing technique, as illustrated in the picture below:

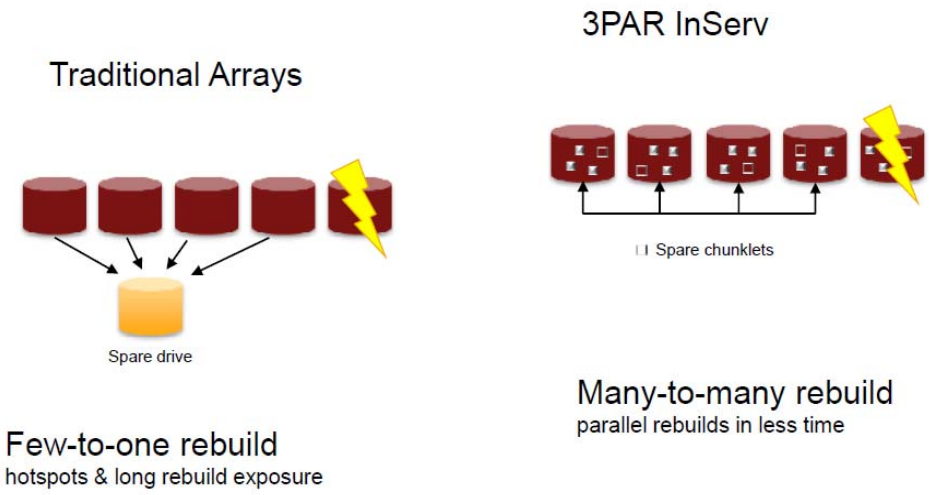

That's why we have experienced a rebuilding time range from 24 hours or longer in the past.

Note: The rebuilding time is less than 2 hours for a 2TB disk per SC8000 or 3PAR, while it took 7 days to rebuild a "RAID" group at the LA Sanitation Agency per one storage vendor's statement.
A:

- Prior to a failed drive replacement, NetApp DATA ONTAP OS will detect a failing drive and copy all good blocks from the failing drive to a global spare drive (Usually, good blocks are at a 98% rate).
- When a copy is completed, the OS will tell a user to replace the failing drive.
- After a new drive is inserted, the rebuilding process will start to rebuild the remaining bad block. Often, the bad block is at 2% rate. Therefore, the rebuilding process is fast.

Note: There is no estimated time obtained for rebuilding a 2TB HDD. Therefore, a filed test is highly recommended to get the best answer.

Q12a: What Do You Think EMC VNX and VMAX Rebuilding Time with a 2TB Disk?

A: Do not know

Q13: Virtualization Concept – Vendor Agnostic vs. NetApp NPIV

VMware is to servers as Dell Compellent is to storage. SC8000 uses NPIV (N-port), which not equal to WWN. NPIV is to VM, a service.

Definition of N_Port: (Source: Wikipedia)

N_Port ID Virtualization (NPIV) is a Fibre Channel facility allowing multiple N_Port IDs to share a single physical N_Port. This allows multiple Fibre Channel initiators to occupy a single physical port, easing hardware requirements in Storage Area Network design, especially where virtual SANs are called for. NPIV is defined by the Technical Committee T11 in the Fibre Channel - Link Services (FC-LS) specification.

May You Tell Me More about This Concept in a NetApp Environment?

A: Skipped due to time running out

Q14: What's the NetApp Vision regarding to a Secondary Storage such as EMC DD?

Note: Gridstore (http://www.gridstore.com/) might be one good technology to address the Disk-to-Disk backup issue at 1/3 cost.

A: NetApp thinks the secondary disk-to-disk storage might not be necessary if enterprise(s) can utilize its highly efficient snapshots to replicate data between two sites or multiple sites.

Q15: Is It Still True regarding the Cisco UCS's Highly Efficient Extended Memory Technology that Can Reduce Memory Costs by Up to 60%?
(Refer to http://media.netapp.com/documents/ds-3299-flexpod-es.pdf)

For example, Cisco UCS servers are much more expensive than HP's. According to an article titled **Service providers trade Cisco UCS for commodity gear**
http://searchdatacenter.techtarget.com/news/2240146725/Service-providers-trade-Cisco-UCS-for-commodity-gear

A quote " One problem is that Cisco UCS costs easily 30% more than comparable servers … gear from vendors like HP…

A: Skipped due to time running out

Q16: Tell Me about NetApp Auto-tiering?

I was told that NetApp started to realize the Auto-tiering is great technology. Therefore, it started to offer this feature recently. Please tell me more about it. See the link below for details:
http://itbloodpressure.com/2012/06/26/take-a-dip-in-the-flash-pool-netapps-auto-tiering/

A: NetApp only does Auto-tiering on SSD and does not believe the current Auto-tiering from most storage vendors is an efficient method because it took too long to move data from third tier to first tier. That's why NetApp creates a different Auto-tiering algorithm. See **NetApp Virtual Storage Tier** - http://www.netapp.com/us/technology/virtual-storage-tier/ for details.

Q17: May You Give Me the Estimate in Comparison with the Table Listed Below?

Compare 1: 90 TB usable, 1 Engine (2 Controllers) Systems		Savings
VMAX 10K	**3PAR V400**	VMAX 10K is 7.3% less expensive than 3PAR
1 Engine (2 Controllers) / 96GB Cache	2 Node (2 Controllers) / 96GB Cache	
100GB FLASH, 300GB FC, 2TB SATA	**100GB FLASH, 300GB FC, 2TB SATA**	
Advanced Software Suite including FAST VP, Virtual Provisoning and TimeFinder	System Reporter, Adaptive Optimization, Virtual Copy, Thin Provisioning	
3 Years HW & SW Maintenance	3 Years HW & SW Maintenance	

Compare 2: 165 TB usable, 2 Engine (4 Controllers) Systems		Savings
VMAX 10K	**3PAR V400**	VMAX 10K is 11.5% less expensive than 3PAR
2 Engines (4 Controllers) / 192GB Cache	4 Node (4 Controllers) / 192GB Cache	
200GB FLASH, 300GB FC, 2TB SATA	**200GB FLASH, 300GB FC, 2TB SATA**	
Advanced Software Suite including FAST VP, Virtual Provisoning and TimeFinder	System Reporter, Adaptive Optimization, Virtual Copy, Thin Provisioning	
3 Years HW & SW Maintenance	3 Years HW & SW Maintenance	

This will be Apples-to-Apples comparison. If possible, can you offer a 3-tier solution quote and compare it against the savings with VMAX 10K solution?

A: Skipped due to time running out

Q18: Advantages and Disadvantages of NAS

The following presentation slide was displayed on the projector screen in Microsoft training class titled MS-20410 "Installing and Configuring Windows Server 2012" as well as in the courseware.

Q19: Tell Me about NetApp Usable Disk Space, as Shown in the Screenshot below:

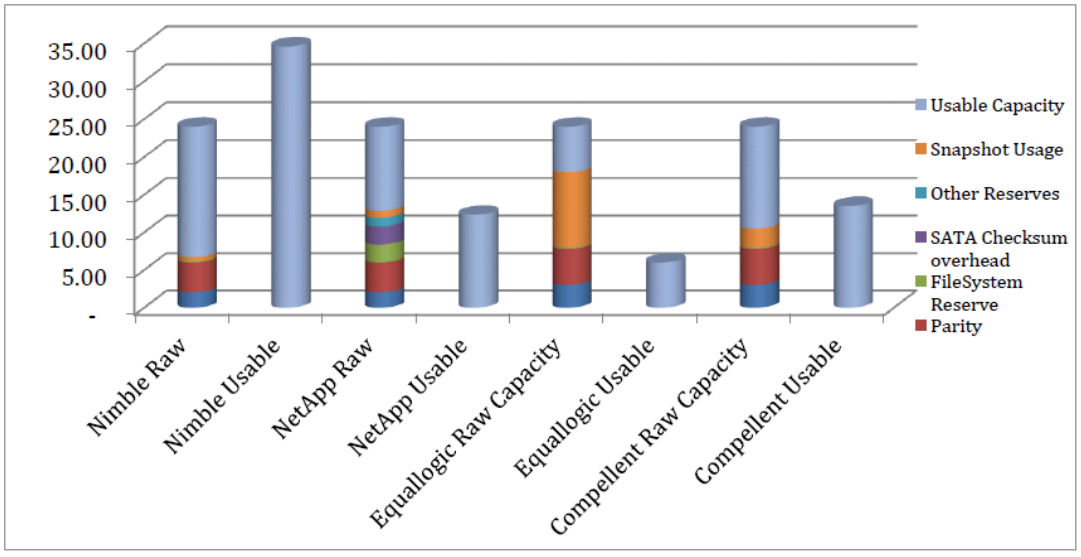

Source: Nimble

Note: Author asked Nimble to supply the detailed info (e.g., which version of DATA ONTAP OS is used in this comparison chart) and will add it into this note once the info is obtained.

A: Skipped due to time running out

Q20: Get the Best Service and Support

When we evaluate enterprise storage solutions, we not only compare the technology, but also look for service. Per EMC, it has 250 to 300 senior engineers just for Los Angeles area.

How many support staff does NetApp have in Los Angeles?

A: Skipped due to time running out

Q21: A Q&A Session Based on a Note Listed in the Link below:

http://www.lacaaea.com/aaeapdf/san/EMC-VNX-5300-VS-Nimble-CS220.pdf

A: Skipped due to time running out

Conclusion

Although NetApp sells its hardware based storages, it claims NetApp is a software company that has about 40 different applications and licenses to be bundled with its storage array systems. The crown jewel features from NetApp are its efficient snapshot with post-compression, thin replication and deduplication in primary storage. The newly released DATA ONTAP 8.1 added the clustering capability that immediately increases its storage capacity with higher availability and fault tolerance that, in turn, attracts more customers for its storage needs.

NetApp can boost performance of an existing storage 20% or more and comes with a "pays for itself in 9 months" guarantee program besides NetApp guarantees customers who will use 50% less storage for virtual environments.

Below is a quote from a report titled "Storage Demand – NetApp Inc." from OTR Global LLC (http://www.otrglobal.com/)

> The one bright spot for NetApp continues to be its FlexPod virtualization architecture through the partnership with **VMware Inc.** and **Cisco Systems Inc.** During FY4Q12, NetApp's FlexPod won business from **EMC Corp.**'s Vblock for a number of sources, continuing a trend OTR Global has tracked for several quarters. A U.S. source said, "FlexPod is still quite popular, and we're seeing more actual FlexPod deals closing. I haven't seen any drop off."

NetApp FlexPod architecture along with its "pays for itself in 9 months" guarantee program garnered its top spot in a fixed period of 2012 from the GOLDMAN SACHS 2012 Storage report and achieved 60% market share in Federal Government.

However, when a final price includes the software licenses with its storage array, an enterprise often finds that the price jumps. Customers know they can get the same storage capacity and similar software features for less money elsewhere. For example, in the SMBs market, many companies are able to find an equivalent solution from Nimble at much cheaper price because Nimble offers snapshots and replication software, including Zero-copy clone free of charges.

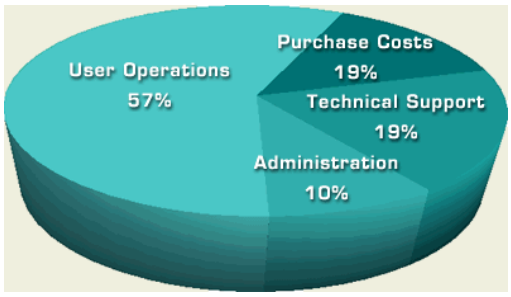

Gartner introduced the idea of "Total Cost of Ownership" (TCO) in 1986, as illustrated in the screenshot from left. The initial acquisition cost is only a part of the equation of TCO. More software purchased, more hidden cost arisen.

Therefore, any company should select a solution in accordance with its own environment needs.

Source: http://www.alligatorsql.com/solutions/tco/index.jsp

Often, CIOs in any organization(s) are being asked to do more with fewer resources, less time and less funding, and transform and guide an enterprise into a lean organization. Therefore, carefully examining its internal IT operations and efficiency, including using available technology in new efficient ways, often leads to trim many hidden costs and yield greater flexibility to focus on core business and pursue the innovation.

Challenge

EMC is still the #1 disk storage market leader. NetApp is no longer a low cost storage vendor if its software is added into the picture. Many customers know they can get the same thing for less money elsewhere. For example, Nimble (http://www.nimblestorage.com/) can compete with NetApp and EMC at the Small Medium Businesses (SMBs) market because Nimble includes all software (e.g., snapshot and replication) in its storage solution for free.

EMC's Isilon is also a big challenge to NetApp in NAS environment (Generally speaking, Isilon is the scale-out storage solutions for media content creation, distribution, and archive workflows.) NetApp also lost many deals to Dell and Nimble due to its higher pricing in SMBs market.

EMC's new VSPEx might be a challenge to FlexPod's success.

Appendix A:

Data Center Fitness/NetApp V-series with Existing Storage Vendor
http://www.netapp.com/us/campaigns/it-efficiency/

Special note: EMC VNX Unified Storage Delivers 25% More Storage Efficiency — Guaranteed - http://www.emc.com/about/news/press/2011/20110308-02.htm
Server Consolidation/Data Center Efficiency
http://media.netapp.com/documents/ds-3299-flexpod-es.pdf

BYOD and VDI
http://www.slideshare.net/NetApp/oak-hills-copy

NetApp wins VMWorld contest
http://media.netapp.com/documents/vmworld-2012-brief.pdf

Everything you've ever wanted to know about NetApp
http://www.netapp.com/us/media/netapp-factsheet.html

Recommended Reading

1. **State of Texas Moves More Than 100,000 State Employees to Microsoft Cloud.**

The State of Texas is moving more than 100,000 employees onto Office 365 at a cost of about $3.50 per user, per month, making it the largest statewide deployment of email and collaboration services in the U.S.

Visit http://www.microsoft.com/en-us/news/Press/2013/Feb13/02-15TexasO365PR.aspx for details.

2. **How New York City is going to Consolidate 50 Data Centers from 40 City Agencies into One Location.**

Visit http://www.informationweek.com/government/state-local/nyc-opens-consolidated-data-center/229219575 for details. (Source: InformationWeek)