Jeremy Li
May 23, 2012

Credit: Dell/ExaGrid/MR2 Technical Lunch Event, Downtown Los Angeles

ExaGrid Systems (Up to 130TB per GRID) - www.exagrid.com

ExaGrid offices are across North America, Europe and Asia. Its product is highly scalable, cost-effective disk backup with deduplication system.

As of today, 4,200+ systems are installed globally with about 1,300 customers. It is application agnostic.

Each ExaGrid's disk backup appliance is comprised of Intel Quad Core XEON processors, enterprise serial advanced technology attachment (SATA) drives, RAID 6 with hot spare, and work seamlessly with the industry's leading backup applications and the appliance typically installs in about one hour.

The disk backup is faster than tape, and restores with greater reliability. ExaGrid. offering either 1Gb or 10Gb Ethernet connectivity with single controller scale out up to 10 maximum, uses its unique landing zone architecture and brings a scalable storage node clustering, also called grid, with post-processing zone-level data deduplication at byte-level, instead of block-level technique to perform a backup directly to disk. This will shorten the backup window in comparison with in-line processes, which will slow backups. Generally speaking, it yields 2-3X faster than Inline and 30-90% faster than tape, per ExaGrid.

On the other hand, all backup data will be transmitted over the corporate network prior to the arrival of the landing zone on its appliance. Therefore, it generates most network traffic in comparison with inline depupe method.

If you have 50TB data with deduplication ratio at 2:1, the 5.8TB disk space will be consumed on the D2D secondary storage, as shown below:

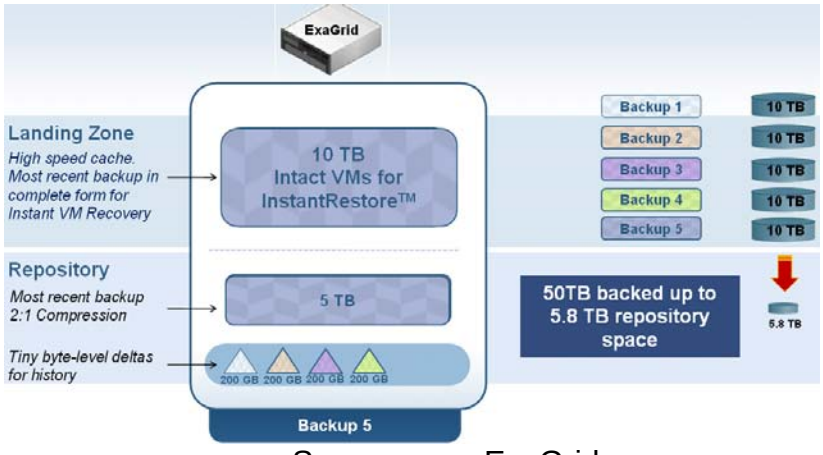

Source: ExaGrid

What level of compression and data deduplication do you achieve? (Source: ExaGrid)

ExaGrid compresses the most recent backup file and keeps it in its entirety. Average compression is about 2 to 1, so a 1TB backup file would be stored as 500GB. All previous backup files are then kept as the byte level changes only, which averages to about 2% of the data, equating to about 20GB for every 1TB of data. This means that if you kept 20 weeks of backup retention, for 1 TB of primary data, the latest backup would be stored as 500GB and the previous backups would be stored as 19 x 20GB (380 GB byte changes) or a total of 500GB + 380GB = 880GB. If you did not use compression and byte-level data deduplication you would need 20TB of storage space. In this example ExaGrid would use **880GB** versus **20TB**. This instance is 23 to 1 data reduction. Overall we see anywhere from 10 to 50:1 reduction in disk consumption.

As of this writing, EX130-GRID is a flagship appliance system from ExaGrid, where 2.40 TB/hour can be achieved in its backup throughput per node. Click on the link http://www.exagrid.com/Products/ExaGrid-Disk-Backup-Product-Line/ for details.

The EX130-GRID system has 320TB RAW disk capacity at its maximum, but 260TB usable disk space is available due to its overhead of RAID-6 with one hot spare in the GRID system.

- Each EX13000E appliance has 16 hard disk drives (HDDs), but 13 HDDs are available for usable disk space. Thus, 26TB usable disk space is available with each 2TB disk drive capacity (13HDDs * 2TB = 26TB).
- 13TB (50% of 26TB) of useable disk space is utilized for its Landing Zone
- The remaining 13TB (50% of 26TB) is utilized for its Repository/Retention Zone [tiny byte-level details for history (snapshots)].
- A full GRID (10) EX13000 NODES has capacity, as shown below:

- 320 TB of RAW Disk Space
- 260TB of Useable Disk Space
- 130TB of Landing Zone
- 130TB of Repository/Retention Zone (providing longer on disk retention)
- Additional CPU (Performance), Memory and throughput per NODE
- Each Node is configured RAID6 (with 2 Parity Drives and a Hot Spare)
- 24TB/HR Throughput

ExaGrid "GRID" Architecture

- Stable backup windows, linear performance as data grows
- Capacity virtualized across nodes – from 1TB to 320TB (RAW) singe GRID (10 EX13000E nodes)
- Deduplication shared across nodes

- Avoids forklift upgrades. Product has no obsolescence
- System right-sized to current data size
- Restores faster from full copy vs. “rehydrating” data
- Lowest TCO
- WAN/MAN replication efficiency is at about 50:1 ratio due to optimized deduplication

Due to its Landing Zone architecture, the capacity planning for a D2D system will be associated with its Landing Zone capacity. (e.g., EX130-GRID has 130TB Landing Zone maximum capacity; therefore, no more than 130TB full backup should be planned). ExaGrid recommends using 100TB maximum capacity for its EX130-GRID system).

However, if a customer has more than 130TB data to be backed up, a customer can use two EX130-GRID systems to accomplish the task, although two EX130-GRID systems will not communicate each other. It really depends on an environment. For example, one EX130-GRID system is used for backing up one share just for the entire messaging system or databases (e.g., Oracle database), while second EX130-GRID system can be used to back up the remaining data.

Please be aware that an ExaGrid system engineer will configure the Grid system for you when a 13TB Land Zone (assume that one-node system is installed) is reaching its full capacity during a non-backup window period in order to let one shared backup folder cross from one node to another node.

Per ExaGrid, all nodes within one GRID system can communicate each other. For example, a messaging backup folder is configured to use both node one and node two. When incoming messaging backup data reaches to node 5, the backup data will be automatically redirected to a Landing Zone configured for node one and node two for its backup.

The Disk-to-disk secondary storage market shares are in the order of:

EMC: Data Domain Deduplication Appliances
ExaGrid: ExaGrid Deduplication Appliances
Quantum: DXi Series Appliances.

Only ExaGrid appliance(s) among top three D2D vendors listed above can avoid a forklift upgrade(s), as shown in the screen shot below:

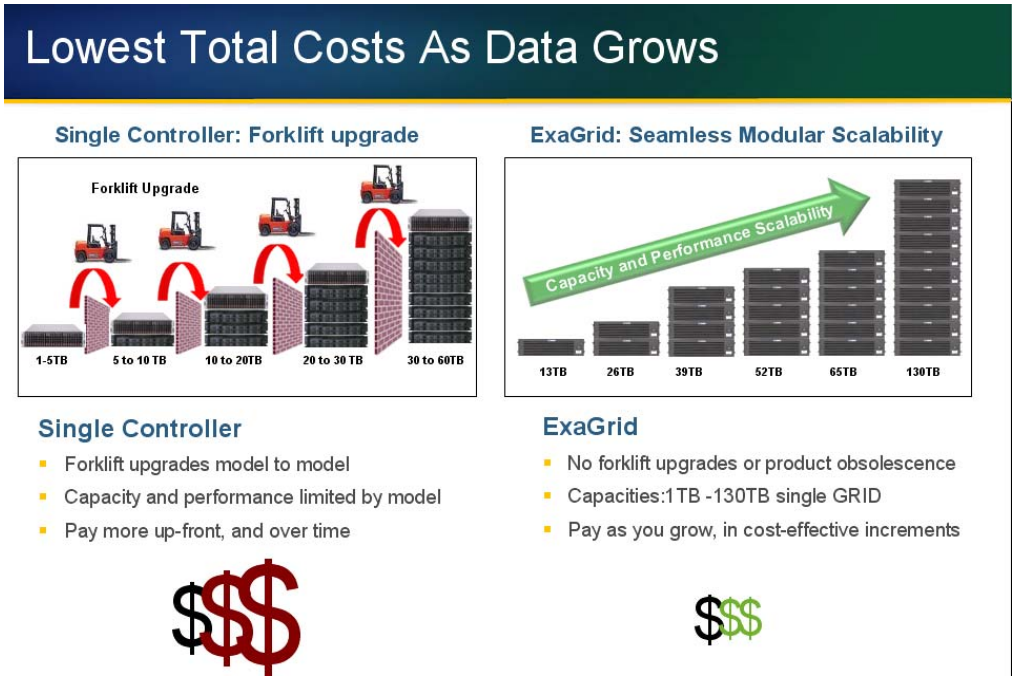

Source: ExaGrid

Each appliance (single controller) is plugged into the switch and virtualized into the GRID, performance is maintained and backup windows will not be increased as data is added.

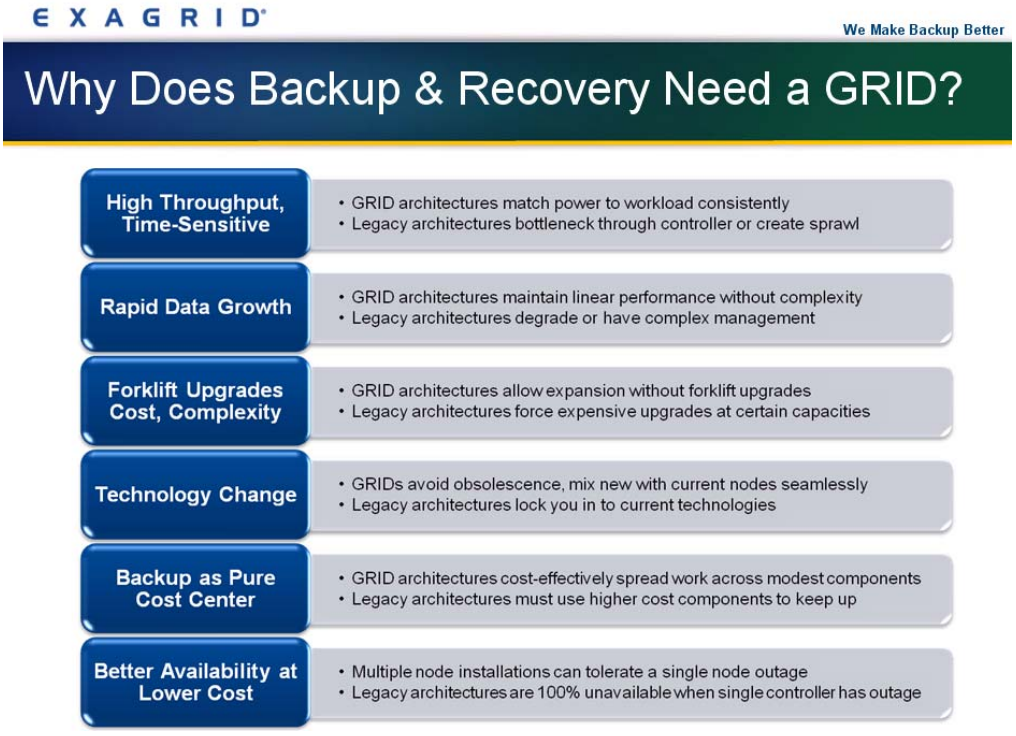

Source: ExaGrid

Generally speaking, ExaGrid appliances use “scale-out” architectures to avoid a forklift upgrade. That is a plus. However, under certain conditions, it might be more cost effective to replace an old, lower capacity appliance with newer and higher capacity appliance to avoid appliances’ sprawl, which will consume more power, energy and footprint at the data center.

Daily Email Report

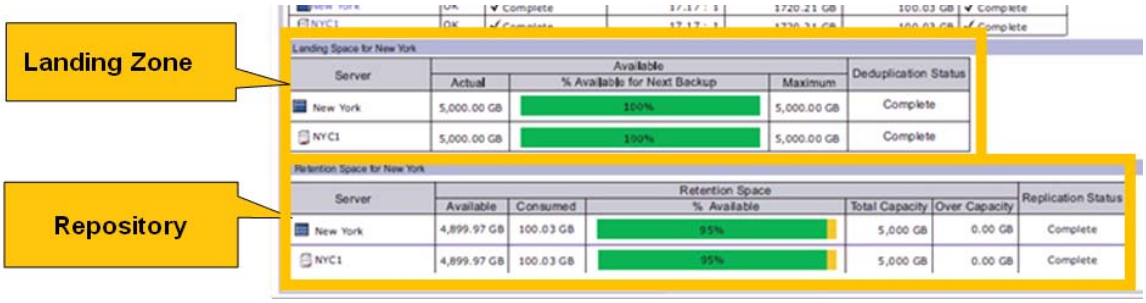

Source: ExaGrid

Management:

All D2D devices can be managed centrally.

Local and Remote Replication

ExaGrid supports both synchronous and asynchronous replication. Generally speaking, synchronous replication is used for Local Area Network (LAN) and Metro Area Network (MAN) for the best data protection (e.g., where no single transaction can be lost in case of a DR), while asynchronous replication is used for a long distance replication (e.g., Los Angeles to New York) for a highest possible performance purpose.
The asynchronous replication is mainly used for disaster recovery (DR) purpose.

In a remote replication scenario, it only replicates deduplicated data to save network bandwidth.

- Please note that a media server must be used at the DR site because the D2D system can not read the backup data and can not present it. It is simply a repository for the data.
- Since the data on the D2D is not in its native format due to deduplication, a landing space for ALL of the rehydrated data is needed. For example, if you have 3 months worth of backup data to be restored (e.g., 100TB raw data, and 10TB after the dedupe), a landing space for ALL 100TB is needed.

Therefore, snapshots are more efficient in a DR scenario. For example, Nimble or NetApp can replicate data to a remote location for DR purpose with highly

efficient redirect-on-write snapshots, and the same data is already in place and ready to be presented immediately. That's a time savings with fewer error-prone.

All software licenses, including remote replication, are included in its appliance, while some other vendors will charge optional software licenses.

Backup Network Traffic

Due to ExaGrid's post-processing zone-level data dedupe architecture, all backup data will be transmitted over the corporate network. Therefore, it generates most network traffic in comparison with inline depupe method.

Logical or Physical Partitions

ExaGrid supports up to two (2) maximum partitions, while EMC Data Domain 600 Series or Quantum DXi 6000 Series can support 100 and 64 partitions, respectively.

NAS (CIFS/NFS) / Symantec OST vs. VLT (Fiber Channel Only)

In the past, VTLs were primarily used for enterprise environments while NAS solutions were used in the midrange arena. New physical networks had to be implemented in FC deployment, while the ISCSI SAN and NAS can utilize on investments already made in typical networking equipment, which uses the Ethernet and IP protocols. However, with 10GbE adoption widely available, the gap for the performance of the backup between the VLT and NAS becomes smaller and smaller.

NAS or iSCSI with 1GbE or 10GbE implementation will always be slower than comparable FC deployments, because NAS deployments communicate on a higher level of the protocol stack (TCP/IP). This makes them inherently slower but has the advantages - a lower implementation cost, and the data and disks actually can be shared between targets (storage) and clients.

ExaGrid supports either 1GbE or 10GbE interface, but does not support fiber channel (FC) due to the reason mentioned above. The Quantum DXi6700 Technical Specifications under "INLINE PERFORMANCE" also demonstrates the trend -

- NAS Interface: 5.0 TB/hour
- OST Interface: 5.8 TB/hour
- VTL Interface: 5.8 TB/hour

At public sector, the working hours are <u>not</u> 24x7 among most agencies, except for law enforcement agencies. Therefore, the overhead of NAS (NFS/CIFS) and Servers can be overcome by using nightly and weekend backup. The savings for

using economical NAS / iSCSI with 10GbE over expensive Fiber Channel is very significant due to FC's complexity with a lot of operating expense (OpEx). Any agencies want less complexity, improved operational efficiency and ultimately, to spend less money.

In addition, Author has been using a private Local Area Network (LAN) - (e.g., 192.168.0.0) since 2003 to backup one public agency's entire network with success by simply enabling second network interface card (NIC) embedded on a server to form a private LAN for its dedicated backup jobs with a $500 Gigabit switch. As a result, all backup traffics do not go through a regular LAN.

Read "Top 3 Reasons Why You Can't Build a Cloud with Fibre Channel" for additional reasons - (Source: CORAID)
http://san.coraid.com/rs/coraid/images/Coraid_Brief_Top3ReasonsWhyYou.pdf?

If a customer is using Symantec NetBackup or Backup Exec, the OST implementation might reduce the backup time significantly. Click on the links below for details:

Data Domain Boost for Symantec NetBackup -
http://www.datadomain.com/pdf/h7143-datadomain-boost-for-netbackup-so.pdf

Five Guidelines as to the Best Ways to Implement Deduplication / NAS is becoming the preferred interface of deduplication targets -
http://www.dcig.com/2011/04/five-guidelines-best-implement-deduplication.html

Caution: Although Symantec NetBackup or Backup Exec does not charge the license fee for the OST implementation, a third party D2D vendor (e.g., EMC Data Domain Boost) might charge the license fee.

In addition, if your media server does not have enough horse power (e.g., HP DL380G6 or Dell R810 or newer), it is not a good idea to deploy the OST in your environment since the OST requires to use high resources on a media server. As a result, the backup performance will be extremely slow.

Conclusion:

ExaGrid is a good D2D storage system and received numerous awards, including Gartner's remark "Top Emerging Vendor" Cool Vendor in Data Protection in March, 2008. ExaGrid's patented byte-level data deduplication technology and most recent backup compression reduces the amount of disk space needed by the deduplication ratios of 10:1 (90 percent) to 50:1 (98 percent).

Due to its Landing Zone architecture, 50% of its entire usable disk space will be used to store the latest native backup data for exchanging its backup and restore

speed. Therefore, the capacity planning for a D2D system must be associated with its Landing Zone capacity.

Today, ExaGrid flagship product is EX13000E, and the EX130-GRID (10 nodes) has 130TB Landing Zone maximum capacity; therefore, no more than 130TB full backup should be planned. ExaGrid recommends using 100TB maximum capacity for its EX130-GRID system.

Here are a few advantages after its Landing Zone capacity penalty:

• Only solution that keeps a backup window from expanding as your data grows
• Seamless scalability from 1-130TB full backup, with no forklift upgrades
• Fastest backups via full servers in a GRID and post-process deduplication
• Fastest restore and instant VM recovery with full recent backup copy on disk
• Lowest price and lowest TCO

Click http://www.exagrid.com/Products/ExaGrid-Disk-Backup-Product-Line/ for Product Overview.

If your capacity planning will surpass 100TB in next three years, you need to look for another vendor's D2D product. But, ExaGrid should announce a higher capacity GRID system, if a single hard disk drive capacity is available from either 3TB or 4TB capacity.

Per ExaGrid, 25% of its customers are using disk capacity between 1TB and 10TB, while the remaining 75% of its customers are using 11TB and more. 50% of its customers are using 50TB and more. That's the market ExaGrid is targeting for.

Most vendors' D2D products look similar and do the same job, but an in-depth look into features and context should match your environment with a right vendor.

If an environment needs a D2D secondary storage system, it is worth the time to evaluate the ExaGrid appliance(s) to exam deeper to find the best price to value for your actual needs to maximize the return on investment (ROI).

Symantec Open Storage Technology (OST) API, can improve the backup performance dramatically. Therefore, the FC block protocol solution for the VTL implement as opposed to CIFS or NFS in NAS environment might not be necessary. Click on the link http://symantec.dcig.com/2009/05/netbackup-ost-delivers-unexpec.html for details.

Any buyer should also consider the follow factors:

- Vendor's reputation

- Total cost of ownership (TCO)
- Resources, including consulting services
- Customers' installed base
- Technical support
- Customers' satisfaction
- Available Technical Training

Last but not least, with 10GbE widely available, a backup solution from economical NAS / iSCSI with 10GbE, or, from OST, will achieve almost the same result with VTL (FC only) solution, but costs much less than FC deployments. In summary, make an award in the best interests of the organization after all factors have been evaluated.

Challenge:

1. **A SAN, NAS or Unified Storage Acting as a D2D Backup System**

Although disk-to-disk backup is now emerging as preeminent backup architecture, many organizations realize that it can be completely eliminated by using its existing storage (SAN, NAS or unified storage) as a secondary D2D backup system to simplify its data center infrastructure by simply utilizing an efficient redirect-on-write (ROW) snapshot with either inline or post process compression technique to write snapshots to high-density and low cost SATA/SAS disks within its primary storage system.

Another advantage is that all backup data will not be transmitted over the corporate network, but kept within its primary storage system. Therefore, it virtually increases the network bandwidth, if a D2D secondary storage system is not deployed in a corporate network. With snapshots replication along with dedupe, corporate data can be saved at a remote site for DR purpose.

For example, most customers from NetApp (www.NetApp.com) or Nimble (www.nimblestore.com) do not rely on a third party vendor's D2D backup system in order to be able to avoid a complexity and additional layer storage infrastructure, thus, greatly reducing the cost of storage in the data center.

2. **Setup a testing or development environment?**

With a unified or converged storage implementation (e.g., NetApp or Nimble), you can easily setup a testing or development environment quickly by utilizing the ROW snapshots, as shown in the screen shot below without consuming a lot of disk spaces. As a matter of fact, the actual disk space consumption under the "Fast cloning" is a fraction of the original disk space.

This approach not only avoids a complicated D2D secondary storage system, but also creates a powerful testing or development environment quickly to reduce the cost and achieve highest return on investment (ROI).

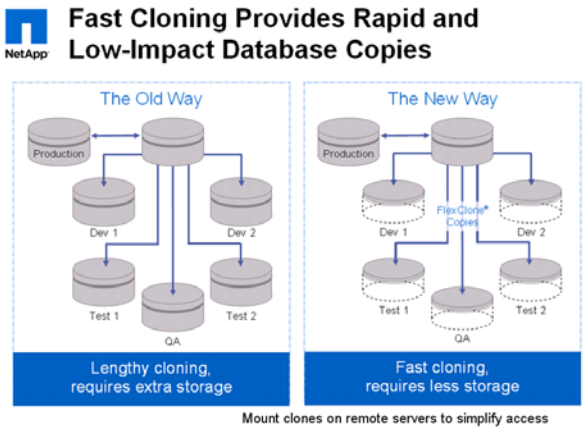

Source: NetApp

3. **ExaGrid does not Support Virtual Tape Library (VTL)**

ExaGrid does not support Virtual Tape Library (VTL), which presents disk based storage systems as tape cartridges, tape drives or tape libraries for use with existing backup application (e.g., Symantec NetBackup or Backup Exec), which thinks a VTL as a physical tape library. The VTL requires fiber channel (FC) interface on the appliance while ExaGrid only supports iSCSI. According to Into Tech (www.infotech.com) analysis, the VTL implementation in a small organization should consider using either FalconStore software or EMC, while Quantum should be considered for Mid-sized Organization.

ExaGrid's position is that iSCSI is much cost effective, while a VTL implementation will add cost significantly because the FC interface and software, which must manage the VTL interface, are required.

An alternative solution over the VTL implementation is to use economical NAS interface or use Symantec Open Storage Technology (OST) API, which will reduce the backup windows dramatically at additional cost. Therefore, the FC block protocol solution for the VTL implement as opposed to CIFS or NFS in NAS environment might not be necessary. Click on the link below for details: http://symantec.dcig.com/2009/05/netbackup-ost-delivers-unexpec.html

Special note:

HP ProLiant G8 server provides two 10GeE NICs via FlexLOM (Flexible LAN on motherboard), as shown in a screen shot below. Any customer can have an option to purchase 10GeE NICs when a new server is ordered.

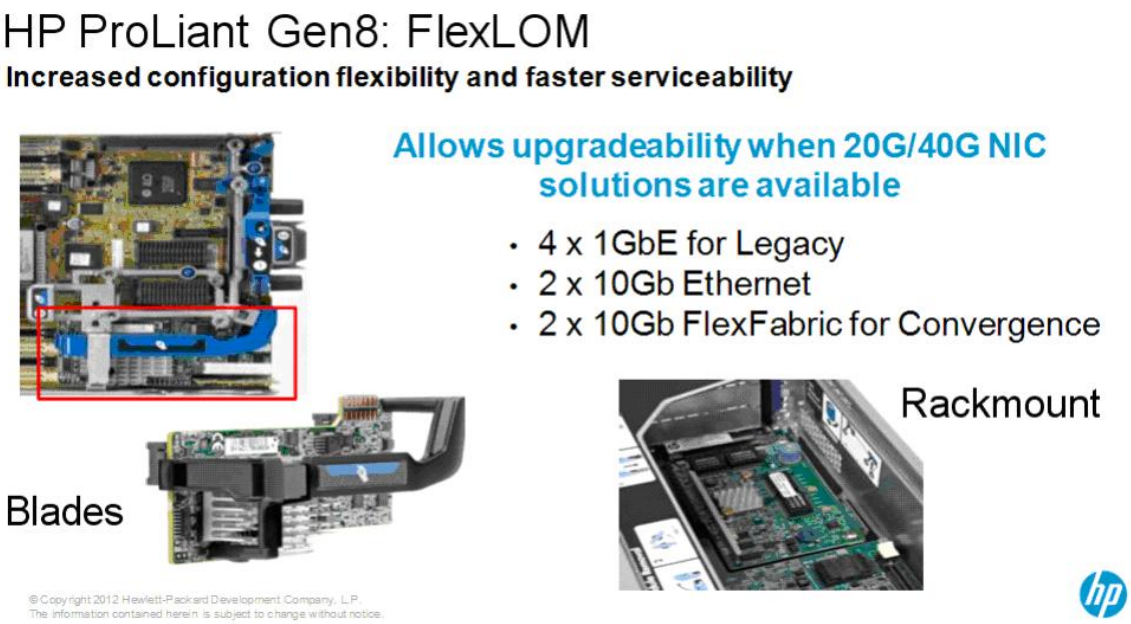

Source: HP

This is another reason the FC block protocol solution for the VTL implement as opposed to CIFS or NFS in NAS environment might not be necessary, in order to reduce the implementation complexity and cost.

4. **Symantec Client-side Dedupelication does not Work with ExaGrid's**

For an obvious reason, a Symantec customer will not get its client-side deduplication, if an ExaGrid appliance is deployed and the Dedupe feature is turned on because both vendors' depupe does not work well together in this field. Click on the link below for details:

http://www.symantec.com/business/support/index?page=content&id=HOWTO23347

Client-side deduplication enables a remote computer that is configured as a Remote Agent for deduplication to send data directly to an OpenStorage device or a deduplication storage folder. Therefore, it not only decreases the storage amount by up to 90%, but also the corporate network traffic. In return, reduce backup windows.

Recommended Reading:

1. Vendor Landscape: Disk Backup (Source: INFO~TECH research group)
   http://www.infotech.com/research/ss/it-vendor-landscape-disk-backup

2. Why EMC Should Be at the Top of Every Enterprise's PBBA Short List
   http://emc.dcig.com/2012/05/why-emc-should-be-top-pbba-short-list.html

3. IDC White Paper – NetApp Storage to Deliver On-Demand Cloud Services Efficiently with High Availability and Seamless Scalability
http://www.emc.com/collateral/analyst-reports/12168w-idc-wp40t-reduced-cs-web-ar.pdf

4. Big Data Backup Environments
http://www.sepaton.com/products/beyond_virtual_tape_library.html

5. Federal Cloud Computing Strategy
http://www.cio.gov/documents/Federal-Cloud-Computing-Strategy.pdf

6. CIO INSIGHT Research Exclusive - 2011 Vendor Value Study
http://circ.ziffdavisenterprise.com/pdfs/Research_CIOI11_0708.pdf