TANEJA GROUP
TECHNOLOGY ANALYSTS

T E C H N O L O G Y B R I E F

Towards Unified Storage

February 2006

THE NET NET

Storage architectures are going through a fundamental change. Whereas we have seen most legacy players continue to offer traditional architectures that present a FC interface towards the host and either FC or SCSI, or more recently, SATA drives on the backend, newer players have developed iSCSI storage using grid principles and made the management of storage much easier. We have also seen players combine NAS and SAN in the same box. Yet others have applied the principles of virtualization to simplify management. We believe all these are major steps towards the concept of Unified Storage that would bring many of these isolated elements together. While a number of traditional players are moving in this direction the most innovative energy is coming from newer, smaller players. One of those that we discuss in detail here is RELDATA that has combined the grid principles with NAS/SAN convergence but focused entirely on a subset of Unified Storage that can be thought of as Unified IP Storage. We believe there is broad applicability for this technology, especially for the medium size enterprises where cost and simplicity are the magic words.

A Step Towards Unified Storage

The goal of unified storage is to effectively utilize all available capacity across heterogeneous storage, allow access to this storage via block or file protocols, simplify allocation and management, and reduce costs of managing storage over time. The technologies that are central to achieving this goal include heterogeneous virtualization and grid computing. Over the past five years or so a number of vendors have innovated along these lines and produced some interesting combinations but no one has yet presented a completely unified storage. We are very close, however, to achieving this objective. We look at an interesting player here that we believe has produced a solid Unified IP Storage offering. But before we look at the details of their architecture let's examine the innovations that have occurred so far in the market. We describe these innovations along five dimensions:

1. **Fibre Channel SANs**

Without question Fibre Channel (FC) SANs have had a significant impact on the way storage is currently deployed and managed. While not perfect, SANs have freed storage from being tied to a single server, enabling it to be easily shared with many different applications and servers across an organization. It also enables simplified storage management via a single console and allows servers to boot from the SAN. However, FC SANs also require expensive specialized components and training, which makes them difficult to implement and manage. As a result, the majority of FC SANs are found in large enterprise datacenters that have access to knowledgeable IT resources and larger budgets. The classic innovators in this category are well known and include

EMC, HDS, HP, IBM and several other smaller vendors like Dot Hill and Engenio.

2. iSCSI SANs

The suppliers of iSCSI SANs can themselves be separated into two distinct categories. The first category consists of legacy storage suppliers such as EMC, HP, and Sun, which have added iSCSI as a separately available front end to their FC storage offerings. The second category of suppliers has demonstrated strong innovation and includes EqualLogic, LeftHand Networks and Intransa. By applying grid computing principles and what we call homogeneous virtualization, to iSCSI, they have done more for making storage easy to use, manage and scale than almost anyone else in the last decade. While providing true storage virtualization, they have focused on the same type of storage across all nodes. It is not surprising that they have been extremely popular in small-to-midsize businesses (SMB) that benefit from the ease of use and lower costs associated with iSCSI SANs.

3. NAS/SAN Convergence

Some storage suppliers have focused on NAS/SAN convergence as a primary differentiator. NetApp is the clear leader in this area. Over the past three years they have added an iSCSI interface that allows them to claim NAS/SAN convergence while retaining the ability to connect to a FC SAN and use their backend storage as FC targets. Otherwise, the remaining system characteristics have not changed and the storage is still packaged with the NAS head so it is purchased as a single entity. And while the backend storage can be partitioned into NAS or SAN, the two environments are still mainly separate. But without question NetApp comes the closest to our definition of unified storage today, even if they are missing some key elements. .

4. Heterogeneous Virtualization

A number of smaller storage vendors have focused their energy on heterogeneous virtualization, i.e. virtualization of storage across disparate vendor offering. This group, led by FalconStor, initially focused on selling virtualization until they realized that the legacy storage vendors were opposed to the placement of a virtualization device in front of their multi-million dollar storage solutions. As a result, these smaller vendors quickly repositioned themselves to either a provider of data protection (FalconStor) or refocused their energy towards the SMB markets. This was a smart move as the benefits of virtualization were well appreciated by the SMB market. They also appreciated the added attention they received from these vendors (traditionally most large players had not focused on the SMB space). Also, while important, performance and scalability were (and are) not as critical for most SMB environments when compared to the needs of large enterprise datacenters.

Other vendors in this category took a slightly different approach. For example, StoreAge utilized an Out-of-Band approach for their device and therefore enjoyed greater scalability and performance. Their heterogeneous storage virtualization solution rivals that of the in-band virtualization vendors and their degree of innovation is just as strong. However, they essentially encountered the same set of issues with

legacy vendors in larger enterprise environments. It is also noteworthy that FalconStor also offers NAS/SAN convergence, including iSCSI/FC and CIFS/NFS. But they do not use the classic grid principles utilized by the iSCSI vendors mentioned above.

5. Innovators Along Traditional Lines

This category of vendors, exemplified by Xiotech and Compellent, utilized the traditional architectural approaches but innovated along the lines of homogeneous virtualization and ease of use, especially in terms of allocation or de-allocation of storage. They have also focused on the SMB markets, which typically are not purchasers of the large legacy storage solutions and are open to purchasing from a smaller vendor that can provide them with individual attention.

Enter Unified Storage

Given the above as a backdrop, we can now define unified storage and explore the expected IT user benefits. We think of unified storage as the amalgamation of many of the trends described above. But at the heart of it is the merger of grid computing and heterogeneous virtualization. In our definition, a unified storage solution would have the following characteristics:

- The ability to communicate equally effectively with incoming data traffic in the form of iSCSI (SAN), FC (SAN) and NFS/CIFS (NAS).

- Support all popular types and forms of storage at the backend. Not only FC, SCSI and SAS/SATA storage but also storage from any vendor that meets popular standards.

- The application of grid computing, which creates a node that consists of a certain amount of compute power, software, storage and network connectivity. An environment may start with a single node and grow theoretically to infinity, but practically perhaps to 32 nodes or so. The additional nodes add both capacity and compute performance, which provides a balanced system as the environment is scaled.

- The ability to be easily managed as if it were a single device from a single vendor. Also there would be no "seam" between the NAS side and the SAN side, except that a certain amount and type of storage would be allocated on-the-fly to one side or the other. And allocation between the two environments would be totally non-disruptive.

- The NAS side can present a single global name space but not necessarily from the use of a clustered file system. Traditional file systems may be used in conjunction with other technologies such that a global namespace can be presented, if desired.

Needless to say this is a tall order. But if we look at the pattern of innovation over the past five years we find that vendors are making strides towards this ultimate goal of a truly unified storage. Of course, as is always the case in this industry, the true innovation seems to come from the smaller vendors or from those larger vendors that have nothing

to lose in attacking a new market. Many of these vendors are mentioned above along with the segment to which they contributed.

Two vendors stand out in terms of their strong drive towards the unified storage goal. FalconStor has innovated along NAS/SAN convergence and heterogeneous virtualization. And NetApp has also taken large strides towards this unified goal with NAS/SAN convergence and implied innovation towards heterogeneous virtualization (V-filer), but has not delivered along those latter lines yet. Of course, with the Spinnaker integration that is anticipated, one would expect a grid-like infrastructure to be implemented as well. To-date, most grid innovation has come from the iSCSI players mentioned above.

Note that regardless of desire or capability, the likelihood of legacy vendors truly enabling heterogeneous storage virtualization is small, since they all have their own storage markets to protect (this includes NetApp). But even homogeneous virtualization from these larger players would be a major step forward, considering that today two storage units of the same model number still behave like two distinct entities from a management perspective.

So who is likely to enable true unified storage?

Enter RELDATA

Little known and just entering the market is RELDATA, a company headquartered in the US and born from the merger of Reliable Data Technology, Inc., and RELDATA Europe, a German company. We now examine the RELDATA solution in terms of its ability to meet the definition of truly unified storage and the likelihood that it will succeed in this very competitive market.

Overview

The RELDATA IP Storage Gateway 9200 is a diskless multi-function data storage and protection management appliance. It enables open storage resource consolidation and virtualization, unified block-level (iSCSI) and file-level (NAS) storage provisioning, and data protection functionality. It allows administrators to consolidate and virtualize popular FC, SCSI, and iSCSI enabled backend disk devices into one single storage volume pool. This pool can be accessed via incoming iSCSI and NFS/CIFS (NAS) protocols over Ethernet. The appliance also supports integrated data protection features including volume snapshots, mirroring and remote replication both on block- and file-level In general, the customer supplies disk systems of his own choice.

IP SAN and NAS Virtualization

By leveraging virtualization, the RELDATA 9200 allows system administrators to assign storage capacity as part of a SAN and/or NAS environment according to the needs of the application, regardless of where the underlying physical storage is located. Rather than managing multiple heterogeneous physical storage devices, virtualization allows the pooling of capacity into virtual volumes that can be allocated on demand. Managing virtual pools of storage greatly simplifies capacity management and planning. It can also eliminate stranded

storage capacity, which is common in non-unified storage environments.

An integrated Logical Volume Manager (LVM) aggregates SAN-level physical volumes and LUNs from storage resource pools to logical volumes and controls the visibility of these volumes to hosts (LUN masking). The LVM provides storage volumes to either an iSCSI target driver for block-data IP SAN provisioning, or to an integrated journaling file system for host-level shared file NAS provisioning (over NFS clients or CIFS mounting points). This allows administrators to present virtualized, scalable storage volumes as iSCSI SAN targets and NAS file systems (CIFS, NFS) to the IP/Ethernet connected Windows, Linux and Unix hosts. This multi-level virtualization approach provides flexible IP SAN/NAS consolidation based on a common storage pool.

Data access control and quota management, particularly important for NAS applications, are supported by policy-based capacity management for Windows, Linux and Unix clients. This includes Microsoft Active Directory support and CIFS client authorization, based on a Primary Domain Controller as well as NFS authorization supported by NIS/Yellow Pages integration.

Snapshots, Mirroring, and Remote Replication

The RELDATA 9200 includes data protection and disaster recovery functionality at no extra charge. Standard features include snapshots, multiple RAID levels, data mirroring, and remote replication along with security, encryption, and compression capabilities.

The RELDATA 9200 supports both copy-on-write and split-mirror snapshots of single volumes and volume groups. Unlike some other snapshot solutions, RELDATA uses a proprietary technology to support an unlimited number of block-level volume snapshots including snapshots of file systems and iSCSI targets. Also, the RELDATA volume snapshots do not cause disk fragmentation and can be read-only and/or read-write. A patent-pending feature allows the administrator to provide snapshots as iSCSI targets, CIFS clients or NFS mounting points for rapid data recovery from SAN or NAS hosts.

RELDATA 9200 supports two types of synchronous data mirroring: mirroring of virtualized volumes and volume groups, including iSCSI targets and file systems, as well as physical block-data mirroring (RAID 1) to leverage attached JBODs.

The RELDATA remote data protection functionality includes uni- and bi-directional iSCSI based block-data and rsync-based file replication over TCP/IP networks, which can be automatically scheduled and used in support of disaster recovery. In addition replicated data traffic can be accelerated using IPSec compression. Remote data transfer fully supports IPSec encryption and CHAP authentication as mandated by the IETF-ratified specification.

High Performance and Availability

The RELDATA 9200 delivers high performance, reliability and availability to

ensure access to data and eliminate potential single points of failure.

The current shipping model includes a motherboard with dual Intel 64-bit 3 GHz Xeon processors and six, 1 GbE ports, where four ports can be aggregated into virtual interface(s) (also called trunking) using IEEE 802.3ad link aggregation. The aggregated ports appear as a single IP address to the hosts, providing increased bandwidth by merging the bandwidth of the individual ports. Two load-balanced Fibre Channel ports and 2 Ultra SCSI ports provide plug-and-play storage connectivity.

In addition it supports volume striping across multiple connected disc arrays. Virtualized volumes of iSCSI target drivers and journaling file systems are provisioned separately and in parallel, which enables a fast direct block-block processing of iSCSI requests.

Reliability and availability features are also included. For example, the RELDATA 9200's firmware, based on a hardened Linux platform, uses just 50MB and is stored on a flashcard, which is more reliable than disk. Also, multiple RELDATA appliances can be deployed in a cluster configuration to provide active-passive failover and load-balanced multi-pathing, which ensures business continuity should an appliance fail. Automatic backup of configurations and metadata to a USB disk drive (stick) enables the appliance to be quickly exchanged in the event of failure and the power supply and fans are redundant to eliminate any single points of failure.

Grid Approach

The RELDATA appliances today are more consistent with the definition of clustering than grid but, according to the company, the upgrade to the grid will be a field upgrade. Today, the services (applications) can run on any of the individual nodes and the nodes are all active-active. If a node dies, the service is started on another available node, with minimal disruption. This is true for both NFS/CIFS and iSCSI traffic. The company intends to add more grid-like features with the use of clustered file system for the NAS side and do the same for the iSCSI side. In the latter case, the failover would then be instantaneous and the nodes would be auto discovered and joined into the group, upon re-entry.

Interoperability

The RELDATA 9200 is designed to operate in almost any storage environment and is vendor agnostic. On the host side, it combines virtualized storage bridging services for iSCSI SAN and CIFS/NFS NAS Ethernet infrastructures. On the back-end it integrates with a wide range of heterogeneous multi-vendor storage hardware including FC, iSCSI, and SCSI enabled RAID arrays, tape libraries, and SAN switches. Software RAID levels 1, 4 and 5 are supported to integrate cost-effective JBODs. Tape library integration supports access control for the robotics controller and the drives. In addition, boot-over-IP-SAN support enables the use of diskless servers and workstations.

Remote data replication operations are executed over TCP/IP networks either on block-level via iSCSI or on file-level via the

rsync protocol. Using two RELDATA 9200 appliances also enables data replication between two FC SANs without the need for iFCP and FCIP gateways. NDMP support facilitates serverless disk-to-tape (D2T) and disk-to-disk-to-tape (D2D2T) local and remote backup and archiving operations.

Note that RELDATA appliance does not support FC on the host side, making it an example of Unified IP Storage appliance.

Configuration and Monitoring

The RELDATA 9200 includes automated management features that make it easy to setup, deploy, manage, and monitor. The appliance includes a micro GUI enabled color front panel display that facilitates appliance configuration and monitoring. In addition, a web based interface enables on-going management and monitoring of storage resource pooling, storage provisioning and data replication. The Java-based management console connects with the appliance using SOAP protocol and is protected by heavily encrypted SSL technology, which allows flexible and secure remote storage resource management.

The configuration of multiple appliances and their combined common data storage is automatically maintained through an XML file configuration format. The entire system autonomously monitors itself and its attached storage to ensure continuous component integrity. Also, iSNS support facilitates the automated registration, discovery, management, and configuration of iSCSI devices on a TCP/IP network.

RELDATA 9200 Usage Models and Benefits

The ability to consolidate heterogeneous SAN and NAS based disk and tape storage environments via a unified gateway appliance has many benefits. In addition, RELDATA includes integrated data protection and disaster recovery functionality, which helps companies solve other important challenges.

Simplified Management and Tiered Storage

Since heterogeneous virtualization is the foundation of this product, virtual pools of storage can be created, based on the performance and availability characteristics of the underlying storage systems. The administrator need not be an expert at each of the storage units that underlie the appliance. Rather s/he can focus on understanding the subtle details of the RELDATA applications. The appliance, in turn, worries about the idiosyncrasies of the hardware/software associated with the disparate physical devices below. The system inherently allows tiers of storage to be created such that the most critical data can be kept on the most performance-oriented storage, whereas less critical (or older) data can be moved to a less expensive storage pool.

IP SAN and NAS Consolidation

The benefits of NAS/SAN consolidation are well known by now. The key here is not that the two environments can work off of the same storage pool but rather that RELDATA has ensured that there is no seam between the two environments. Volumes on either

side can be expanded non-disruptively; they can also be shrunk non-disruptively and the excess released back into the virtual pool. These may look like simple tasks but the current crop of NAS/SAN converged products in the market does not necessarily support such seamless operations.

Data Protection and Disaster Recovery
All the basic data protection applications are integrated into the appliance. The key is not that these applications exist but rather that they are tightly integrated into the environment. There is no upper limit on the number of snapshots a volume can create; the snapshots can also be used to create a disaster recovery scenario whereby the snapshots are taken periodically and replicated to a remote site, without impacting the primary application. The RELDATA appliance comes standard with this suite of data protection applications at no additional charge.

Taneja Group Opinion

There is a reason unified storage has not happened yet. It is hard to do. As we have seen in above discussion, vendors have taken strides towards solving a variety of challenges underlying unified storage. Some attacked the grid challenges; others focused on heterogeneity, yet others on NAS/SAN convergence. It has also been clear to us that the legacy vendors have made incremental improvements but have simultaneously tried to protect their existing business. That has blocked them from true innovation. The smaller players have nothing to lose and therefore we have seen them innovate along many of the dimensions needed for unified storage. In our view, the key to designing a unified storage is to develop the unified architecture from the very outset. If you don't, you end up with some aspect of the system to be sub-optimal. This is the advantage RELDATA enjoys. But even for RELDATA, they needed to take on a finite challenge and they did this by dropping support for FC on the host side. For the market they are focusing on, we believe this is the right tradeoff to make. Their product performance and software suite are enterprise caliber and we like the fact that all basic data protection is built into the platform. While we think they could (and eventually probably will) go into the larger enterprises, medium size companies, whose storage capacity and management needs are not dissimilar to those of larger enterprises but they need fuller integration and greater value for the money, represent the sweet spot. These companies are also more likely to have a strategic relationship with a value added reseller that RELDATA is targeting.

The big storage suppliers, at least to-date, have been focused primarily on enterprise data centers, which need their portfolio of high-end storage devices, specialized add-on software, and extensive service offerings. As unified storage continues to evolve, we expect the big storage suppliers, led by NetApp, will begin offering unified storage solutions over the next 12-18 months. However, this will be a challenge for these vendors since their significant installed base of legacy systems must be included, which will undoubtedly lengthen their time to market.

In the meantime, RELDATA is well positioned to serve the midsized business market. Our research in the US midsized market supports the need for an easy to use, low cost appliance that unifies heterogeneous SAN and NAS environments, is vendor agnostic, and includes advanced mirroring, snapshot, remote replication, and virtualization functionality.

Tiered storage is practically non existent in the midsize companies. RELDATA can easily present its value proposition as one of reusing existing storage for tier 2 and 3 activity, while allowing the customer to choose the primary storage it wants to buy. The customer would have a hard time finding their existing vendor from delivering on such promises. Having said that, it is clear to us that the major players have understood the message that the midsize enterprise space is important for their future success. They are actively working to move towards the unified storage concept, either on their own or via an acquisition. So time is of the essence.

RELDATA's success will depend materially upon its go-to-market strategy. The primary storage market is cluttered with so many players that it is often hard to find differentiation. That raises the importance of the go-to-market strategy for RELDATA. As a new player they have to attract the attention of the US VARs that are being wooed by every storage vendor these days. But RELDATA does seem to have a product advantage as VARs look for something that gives them an edge.

The biggest competition for RELDATA will come from the new crop of iSCSI players. Most have focused on the SMB space, their iSCSI appliances are easy to use and deploy. But they do not typically have an architecturally elegant way of delivering NAS/SAN convergence. And typically they charge extra for data protection. Of course, NetApp is preparing for a big fight in this arena with their upcoming Spinnaker-technology-based unified storage product. So timing is crucial for RELDATA to get a foothold in the market.

The RELDATA 9200 appliance addresses many key issues faced by storage administrators and includes data protection features at no extra charge. We are encouraged by the architectural approach they have taken but, as we have seen over and over again, the road to success requires much more than technology.

Based on what we have seen so far, however, and the promise of new features to come we think RELDATA is worthy of consideration, especially by the midsize enterprises.

RELDATA Gateway Connections