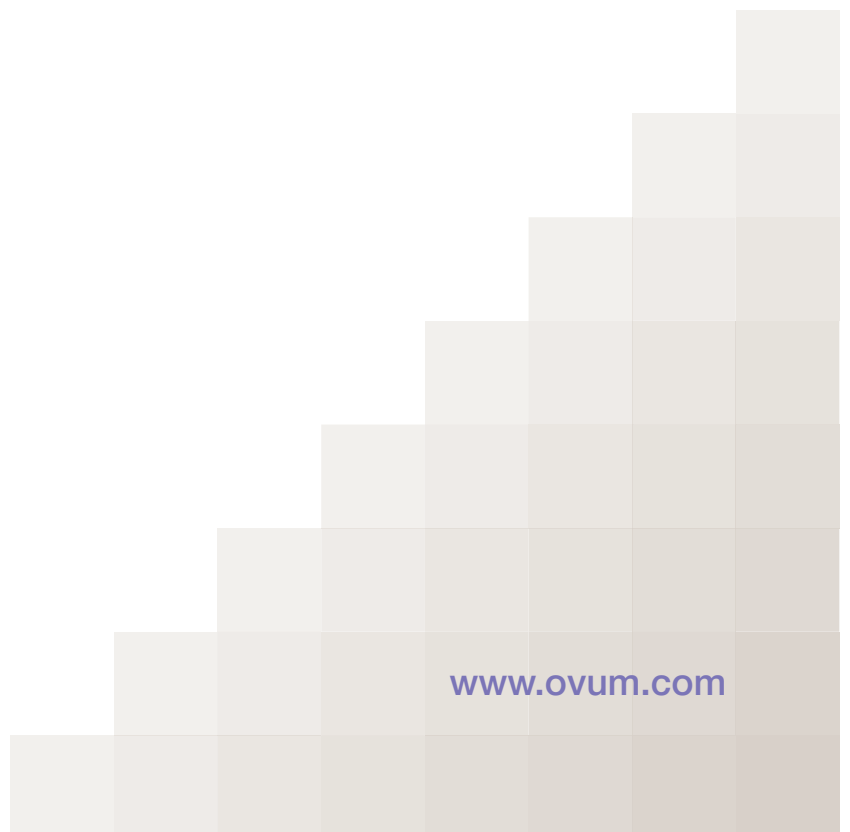




Storage functionality – the intelligent controller versus the network

16 February 2006

Carl Greiner



www.ovum.com



Table of Contents.....	1
Storage functionality – the intelligent controller versus the network.....	2
The current storage landscape and trends.....	2
The storage challenge: simplify the infrastructure.....	4
Storage virtualisation – the enabler.....	5
Network versus storage array controller-based alternatives.....	6
Simplification: an IT imperative.....	8



Storage functionality – the intelligent controller versus the network

There has been intense industry debate as to where critical storage functionality should be placed to optimise and leverage the capability across heterogeneous environments. While storage controller-based functionality remains the most common implementation today, some in the industry tend to give preference to placing more ubiquitous functionality into the network, either within the switch/director itself or as a networked appliance. The most common rationale for network-based functionality is that it would provide a more open, 'leverageable' environment, while minimising vendor lock-in. We believe this is not entirely true, since whatever functionality or solution a data centre implements, it will tie them into that vendor's particular code. With this in mind, what will future high-end storage controllers look like, and where will heterogeneous functionality reside?

The current storage landscape and trends

Tiered storage solutions

As storage capacity requirements continue to grow at a 50%-plus blended rate (with capacity solutions at 80%-plus and modular/enterprise fibre channel (FC) solutions exceeding 45% and 30% respectively), the supporting storage infrastructure continues to challenge most IT organisations (ITOs). Storage management and automatic provisioning functionality is just beginning to mature. In addition, slower large-capacity serial ATA (SATA) or high-capacity FC drives are continuing, and will continue, to be the fastest-growing storage tier by a two-to-one rate compared to the fast FC drives. ITOs are realising that slower capacity drives (300–400Gb) will meet all but the most stringent transaction and database requirements. It is interesting to note that the large capacity drives, with their longer rebuild times and concurrent impact on application performance, are not yet driving a widespread move to RAID 6. In fact, some storage companies do not offer this as an option. This is certainly a trend worth watching.

NAS and SAN infrastructures are converging

Most large global data centres, and an increasing number of small and medium-sized enterprises (SMEs), have made major investments in storage area network (SAN) infrastructures, and are beginning to merge their SAN and network attached storage (NAS) using a NAS gateway to link into the SAN. Fibre channel continues to increase in speed, with 2Gbit/s now common and 4, 8 and 16Gbit/s speeds on the horizon. In addition, 1Gbit/s Ethernet is now a data centre network staple, with 10Gbit/s becoming more affordable. To date, iSCSI implementations have primarily



been in distributed line-of-business (LOB) implementations and mid-sized enterprises or data centres. However, this technology is beginning to see some enterprise data centre traction. While these technologies are often viewed as 'or' technologies, we believe that they will be 'and' technologies for at least the next ten years. As this mixture begins to permeate the data centres, we believe by 2008 best-practice ITOs will merge the current network infrastructure groups with the server and storage I/O interconnect specialists to effectively meet data centre communication requirements.

Storage virtualisation: a requirement for more effective storage management and automation

Storage virtualisation is an emerging technology that creates logical abstractions of physical storage assets and promises to simplify storage administration while increasing utilisation (from 40–50% to 70–80%). While all major vendors and many smaller corporations offer solutions (both block and file based), all use a variety of implementations (in-band or out-of-band – appliance, server, storage controller, intelligent switch/director-based), and each has their pros and cons. While we ultimately believe both the intelligent storage controller and switch approach will prevail, it is important to realise that storage virtualisation has no standard measure of definition or interoperability by an industry group or a standards body. As a result, everyone has a different impression of what it is and what it provides to end users.

Storage and data lifecycle management driven by ITO policies is a journey and normally implemented one application at a time

The unfortunate reality is that information/data lifecycle management is 90% process and 10% tools/automation. Not only do business processes and data requirements have to be understood, but also information classification exercises (creating storage classes and mapping applications and pooled storage devices to the appropriate class) must be completed if data lifecycle management (DLM) is to become a reality. ITOs have to understand their own operational processes and provide a solutions framework to the business users, giving them guidance as to how classifications should be carried out. Key to instilling reality into the exercise is the inclusion of cost metrics for the solution(s) chosen for particular applications/data elements. In conjunction with virtualisation, DLM is an enabler for dynamic policy-based storage management and provisioning. Having incremental success points plus a supportive line of business requiring this function and willing to participate in the classification effort are key enablers.

Stringent regulation requirements are being addressed by capacity disks

Regulation and compliance obligations are requiring that certain data be made available in minutes to low single-digit hours. Automated tape solutions can meet such requirements provided the information is still resident within the library.



However, with SATA/FC capacity drives becoming more cost effective, many enterprises are welcoming a disk-based compliance solution and are happy to push their tape-based alternatives further down the storage hierarchy targeting long-term archival requirements. It is interesting to note that capacity disk solutions are also encroaching on the archival space, as back-ups, compliance and archival requirements are beginning to blend into a total solution based on the lifecycle of data. We believe tape will remain the primary off-site back-up and archival method, with capacity disk being used for the most stringent requirements.

The storage challenge: simplify the infrastructure

Complexity continues to be a key inhibitor to ITO and business unit competitiveness, and it is only through a simplified infrastructure and process optimisation that the free (though secure) flow of information can be realised to address today's demanding business requirements. The ability to simply and easily share and manage data requires integrating IT and business processes, storage assets optimisation and the elimination of silo organisations/domains.

Critical to simplification of the storage infrastructure is the requirement to unify and consolidate the storage environment. This trend has been under way for the past few years as NAS and SAN environments have been merging and virtualisation has begun to address increased utilisation of extensive heterogeneous storage assets. We believe the ability to mix and match disk technologies and architectures (mid-range and enterprise, possibly from different vendors) as required is also crucial to ongoing simplification, with the selection determined and managed by user-defined, service-level policies. This requirement would take the current NAS/SAN consolidation a step further by separating the storage controller personality from the physical disk storage.

Conceptually, an independent storage controller could include an enterprise head, a NAS head (gateway), a modular head, and potentially a tape/tape emulation head. The physical storage, which is common to all storage incarnations, could then be separated and treated as a pool. This goes a long way towards simplifying the storage environment and enhancing asset utilisation, by providing a flexible, adaptable infrastructure. The most effective implementation would be where the personalities of both the enterprise (scale up) and modular (scale out) solutions are contained within common hardware and control code, making the differences moot. Both Hitachi Data Systems and Network Appliance are closest to providing this type of simplification, since both companies have common hardware and support code between their respective storage array personalities. In addition, virtualisation provides a single pool of tiered storage under common management – one set of tools, one set of processes. This further simplifies the infrastructure and begins to break down the technology silos inherent within most data centres.

There are two key functions required to make such a controller a reality. The more obvious is the virtualisation engine's capabilities, and the second is the architecture



of the controller itself. Like servers as opposed to PCs, or directors as opposed to switches, the independent controller must be a real storage controller and not some other device emulating controller functions. Storage controllers are designed for high availability and to maintain state and cache coherency. Physically separating the controller from the storage does not eliminate any responsibilities of the controller, if anything, it makes it even more important.

Storage virtualisation – the enabler

Storage virtualisation is technology implemented in hardware and/or software that masks the underlying technical complexity and incompatibility of disparate storage systems with a cohesive, unified and heterogeneously accessible virtual storage platform. This enables organisations to group their various investments in storage as a central cohesive pool of storage resources, and be efficiently allocated and simultaneously accessed by heterogeneous computing platforms. In effect, it is a technique used throughout IT environments to simplify what tends to be a relatively complex underlying infrastructure. The idea is to separate the logical presentation of a resource from its physical implementation – presenting a simplified, seamless virtual view of the resource to applications and administrators. Storage virtualisation helps the storage administrators perform such tasks as creating storage tiers, and doing back-ups, archiving and recovery more easily, and in less time, by disguising the actual complexity of the SAN infrastructure.

Storage virtualisation allows:

- increased resource utilisation by combining the storage capacity of multiple heterogeneous arrays and architectures into a single pool
- improved personnel productivity by enabling administrators to manage the pool of storage from a single user interface
- the efficient creation of storage tiers, which enables the matching of actual storage types to the value and requirements of the data
- introduction of common storage management and copy/replication services
- consistent automation of resource-intensive operations (such as provisioning, performance and capacity management), and makes data lifecycle management initiatives achievable.

Early implementations of storage virtualisation have generally centred on infrastructure change requirements. A common justification has been to effectively move data from one array to another, either to facilitate an upgrade of existing hardware or to change vendors. Another centres on consolidation, focusing on process optimisation efforts (silo elimination), rationalisation (fewer different types of similar products) and vendor reduction. Virtualisation in these cases can minimise the technical issues associated with multi-vendor environments and the beginning of the rationalisation of multiple multi-path device drivers, multiple replication services and different management interfaces, while optimising asset utilisation.



Given the industry trends, technology enhancements and ongoing ITO challenges, the primary question emerges as to where advanced storage functionality will ultimately reside and how solutions will evolve.

Network versus storage array controller-based alternatives

Where the intelligence to simplify the storage infrastructure and overall environment resides is an interesting question from a technical standpoint, but of less interest to the data centre operations manager who is trying to deliver consistent service to the users. The transparency that storage virtualisation provides for storage assets should also apply to the storage virtualisation solution itself. However, without standards this will prove to be problematic. Just as it remains difficult to intermix different vendor SAN directors/switches and maintain unique functionality, so too will be the ability to intermix virtualisation solutions. The abstraction layer that masks physical from logical storage can reside on a server (host-based) within the storage network as an appliance, or be integrated with the SAN switches (network-based) or in the storage array controller (array-based).

In addition, there are different methods for implementing the virtual storage transport, including in-band (in the data path, both data and control) and out-of-band (separate paths for data and control). Again, the implementation is interesting technically, and in effect becomes almost a religious issue as to what the superior elements and the faults of the different approaches are. The truth is that most implementations are hybrids and tend to blur the technical implementations. As a result, we believe the determining factor of where the intelligence will end up will be installation-specific and normally tied to solving tactical issues with incumbent storage vendors. This is unfortunate since the decision is extremely strategic and the integration of the approaches will be particularly difficult over the next three to five years.

Network-based – intelligent switch, appliance

So far the acceptance of the intelligent switch (from Cisco, Brocade and Mc Data) has been lacklustre and has not gained much momentum. The primary reason is that every vendor's solution supports different switch-based hardware for the functionality engine, and there are no emerging standards to address this issue. As a result, the supported functions tend to be one-offs and, to date, the required hardware has been extremely expensive, thus difficult to cost justify. Further complicating switch-based solutions is the requirement for two function engines (one in each switch/director) to address high availability and eliminate a single point of failure.

Network-based appliances have seen more success (for example, IBM and Veritas/Symantec – both can also be switch based), but have been focused on particular data centre applications, a single domain (for example, Windows), or



smaller data centres. These appliances also require a clustering mechanism to provide high availability. We believe that over the next two to three years we will see more appliance-type solutions embrace/migrate to switch/director implementations and costs will become more competitive. Both of these approaches have to address the issue of maintaining state and/or cache coherency. This adds a significant level of complexity to the design, and one that gets more difficult to solve as systems scale up/out.

Intelligent storage array controller

By design, storage controllers have had to deal with state and cache coherency from the beginning, as data is stateful and persistent. However, the challenge these days is in physically separating the controller from the commodity disk media. Currently, the most robust solution on the market for virtualisation, heterogeneous support and scalability (up to 32 petabytes) is the Universal Storage Platform (USP) from Hitachi Data Systems (also re-marketed by Sun and HP). A miniaturised version, the NSC55 was announced in July 2005 and is aimed at the needs of the SME market; it offers all the same features and functions of the USP. This intelligent control unit has normal storage subsystem redundancy and performance, plus a crossbar switch ASIC. It creates a high-performance switched fabric that supports multiple, simultaneous point-to-point communications between heterogeneous storage network devices, disk arrays and intelligent storage management software. It is interesting to note that it is possible to install a USP with as little as 375Gb of tier-one internal capacity – effectively a standalone storage control unit – in front of a heterogeneous SAN, NAS and mainframe storage infrastructure, enabling a cohesive virtual storage environment with existing storage assets. The control unit also supports logical partitioning for additional flexibility in addressing performance objectives and security.

As with most other virtualisation implementations, native functionality associated with other vendor arrays being virtualised is nullified. However, the virtualisation layer introduces common management and control unit functionality for all externally attached heterogeneous storage. This allows the lower-cost tier two and tier three storage connected to the USP to have the same functionality as tier one – including such USP features as storage-agnostic, heterogeneous replication, logical partitioning of disk, cache, ports for the creation of private virtual storage machines, and software for regulatory compliance. The USP and NSC also enable customers to host less critical mainframe data on their cost-effective, tier-three SATA storage systems, while utilising their high-end systems for critical operational data. Such solutions take customers beyond a conventional enterprise storage system, leveraging common software capabilities and processes beyond a single box – making the solution truly heterogeneous. HDS has established the standard for storage array controller-based virtualisation solutions, and we believe it will be the one against which most other solutions are measured.



Simplification: an IT imperative

Dynamic business requirements continue to drive 50%-plus annual storage capacity growth rates, making the supporting storage infrastructure a major item of focus for ITOs. Simplification through process optimisation and automation is an IT imperative. Storage virtualisation is a primary enabler, allowing enterprises of all sizes to begin to simplify their storage infrastructures and address data lifecycle requirements. Virtualisation and the introduction of common management and functionality across heterogeneous arrays provides an opportunity to introduce dramatic efficiencies to storage management operations, resulting in higher storage administrator efficiencies, lower costs and improved overall IT responsiveness to business requirements.

Ultimately, storage decisions will no longer be about modular versus enterprise array, or fibre channel versus SATA/high-capacity FC drives – they will be more reflective of business value. We believe that storage array, controller-based and network (intelligent switch)-based solutions will gain the most traction at the point of virtualisation. However, storage array controller-based solutions will provide the most robust and functionally rich implementations. Successful IT organisations will put infrastructure virtualisation on the short list of things required, and storage will be one of the early domains to be addressed due to its dynamic growth and sensitivity.

Ovum does not endorse companies or their products. Ovum operates under an Independence Charter. For full details please see www.ovum.com/about/charter.asp.

For full details of Ovum's citation policy, see www.ovum.com/media/citation.asp.

Whilst every care is taken to ensure the accuracy of the information contained in this material, the facts, estimates and opinions stated are based on information and sources which, while we believe them to be reliable, are not guaranteed. In particular, it should not be relied upon as the sole source of reference in relation to the subject matter. No liability can be accepted by Ovum Limited, its directors or employees for any loss occasioned to any person or entity acting or failing to act as a result of anything contained in or omitted from the content of this material, or our conclusions as stated. The findings are Ovum's current opinions; they are subject to change without notice. Ovum has no obligation to update or amend the research or to let anyone know if our opinions change materially.