

Storage strategies - keeping it simple

Clive Longbottom, Service Director

Quocirca Comment

There is a perception that storage is now relatively simple. Although data volumes are rising rapidly, hardware costs continue to tumble, and any problem can be solved by just throwing more disks at it.

Unfortunately, this misses the point; the business impact of a poor storage strategy has to be considered. For example, searching through 1 terabyte (1 thousand gigabytes) of data is a lengthy process: if that data can first be "cleansed" through the use of deduplication technologies to only be, say, 500 gigabytes, the speed of response to business queries will be far faster – and as the results are based on a better quality of data, they will also be more accurate and useful.

That said, even at the hardware level, there have been changes that are requiring organisations to revisit their storage approaches. In the past, where applications were installed on single physical servers, there was a need to ensure that data exchanges between various applications was possible – and meant a lot of back-end data transfers going on which used up large amounts of core local area network (LAN) bandwidth. The solution to this was to provide storage area networks (SANs), based on a separate network using fibre channel to provide high speed data transfer while leaving the LAN available for user traffic.

Virtualisation is also beginning to change the approach many organisations take to storage. Networks can now be virtualised and back end data transfers can be provided with their own high speed LAN partitions, based on 1Gb or 10Gb cores, with 40Gb expected soon. By using fibre channel instructions carried over standard Ethernet (FCoE), existing storage subsystems can still be utilised, but without the overhead of running a specific physical network for them. Moreover, different storage architectures are either appearing – or making a comeback. Blade-based computing or high-density highly engineered rack mounted servers are seeing direct attached storage (DAS) and network

attached storage (NAS) systems becoming more commonplace again. DAS, once the scourge of the data centre due to high rates of disk failure, can now be virtualised to be a peer medium along with NAS and SAN systems.

"Just a bunch of disk" (JBOD) and massive arrays of idle disks (MAIDs) are being implemented alongside highly intelligent self-contained storage systems to provide fault tolerant systems, with the virtualised pool of storage being partitioned as required to provide redundant arrays of (virtualised) inexpensive disks (RAID) that can meet the needs of many workloads in the data centre.

Serial attached SCSI (SAS) and serial advanced technology attachment (SATA) drives are now the dominant connection technologies in use. SCSI-based drives have historically been the choice of server manufacturers due to more complex in-built command capabilities, higher disk speeds and lifecycles, but SATA has been evolving fast and is now a highly competitive and cost effective alternative to SAS.

"Workloads" rapidly become the focus when any organisation is looking at its storage strategy. Whereas file and print storage needs have not changed much over the years, the quantity and speed of transactional data have grown enormously. The problem is that the laws of physics continue to apply; it is only possible to read information from a spinning magnetic disk at a certain rate. The disk read head has to be moved and the disk itself spun to get the data to just the right place for the head to read it. The data then has to be moved from the disk to the server processor itself over some form of transport. The transport mechanism has been evolving, from DAS-based connectors through fibre channel to FCoE, iSCSI and Infiniband, each of which has its own pros and cons. The constraining factor now, however, is often the disk itself – new technology is coming to market to meet the need for high speed data access. The most important is based loosely on the same technology as is used in consumer equipment,

such as flash memory in cameras - solid state disks (SSDs). They have no moving parts and can provide data at extremely high speeds to a data transport. There are downsides - SSD is based on data cells, and each cell can only be written and erased a certain number of times before it becomes unusable. Vendors build in intelligent algorithms to their SSD-based systems to ensure that cells are not continuously hit, so ensuring that SSDs have an optimum working life.

Another approach is silicon-based dynamic random access memory (DRAM), effectively using server memory for temporary storage - used only for the most demanding of data access scenarios.

DRAM is far faster than even SSD and is held close to the server CPU with high speed, short distance data buses, giving very high speed of data access. However, with basic disk drives now costing pennies per gigabyte of storage, DRAM is still several orders of magnitude more expensive.

Right at the back end of data storage requirements is long term archival. Historically this has been served through the use of tape and although there have been many predictions of the death of tape as a storage medium, it still has its part to play. With data retention laws becoming more aggressive, long term storage is a major issue for many, and disk-to-disk archival is really only useful for short term storage. Therefore, Quocirca expects to still see linear tape-out (LTO) being used for many years to come as part of a well architected overall storage strategy.

So, what approach should be taken? Quocirca recommends the following steps:

- Cleanse data through using data deduplication
- Prioritise data storage based on workload requirements
- Use in-memory storage as a "tier 0" if your business can warrant the expense
- Use SSD as tier 1 storage for the majority of low-latency data access needs
- Use SAS/SATA based storage systems as main tier 2 storage, optimising the lifecycle of existing storage systems and adding JBOD/MAID as required
- Cascade older SAS/SATA systems as tier 3, file and print based storage units
- Use LTO tape systems for long term archival

Such an approach combined with suitable storage virtualisation that is aware of how different storage assets can meet different workload needs will provide a long term strategy that should enable organisations to optimise the manner in which it deals with existing data, but also to change as the market dictates.

[This article first appeared through Ingram Micro](#)

About Quocirca

Quocirca is a primary research and analysis company specialising in the business impact of information technology and communications (ITC). With world-wide, native language reach, Quocirca provides in-depth insights into the views of buyers and influencers in large, mid-sized and small organisations. Its analyst team is made up of real-world practitioners with first-hand experience of ITC delivery who continuously research and track the industry and its real usage in the markets.

Through researching perceptions, Quocirca uncovers the real hurdles to technology adoption – the personal and political aspects of an organisation's environment and the pressures of the need for demonstrable business value in any implementation. This capability to uncover and report back on the end-user perceptions in the market enables Quocirca to advise on the realities of technology adoption, not the promises.

Quocirca research is always pragmatic, business orientated and conducted in the context of the bigger picture. ITC has the ability to transform businesses and the processes that drive them, but often fails to do so. Quocirca's mission is to help organisations improve their success rate in process enablement through better levels of understanding and the adoption of the correct technologies at the correct time.

Quocirca has a pro-active primary research programme, regularly surveying users, purchasers and resellers of ITC products and services on emerging, evolving and maturing technologies. Over time, Quocirca has built a picture of long term investment trends, providing invaluable information for the whole of the ITC community.

Quocirca works with global and local providers of ITC products and services to help them deliver on the promise that ITC holds for business. Quocirca's clients include Oracle, Microsoft, IBM, O2, T-Mobile, HP, Xerox, EMC, Symantec and Cisco, along with other large and medium sized vendors, service providers and more specialist firms.

Full access to all of Quocirca's public output (reports, articles, presentations, blogs and videos) can be made at <http://www.quocirca.com>