Storage

Storage

Introduction

If you are thinking about a server project then just as you are probably thinking about server virtualisation, you should also be thinking about your storage options. This Technology Unscrambled Guide intends to highlight the options available and some of the definitions around storage – to assist in that decision.

The adoption of Storage Area Networks (SAN) is on the increase and often goes hand in hand with Server Virtualisation.

As the demands placed on technology grows – and the reliance on company systems and server infrastructure inevitably increases, so too does the reliance on the availability of the information stored. Hence, decisions regarding storage are an integral and significant part of any server project / decision.

If you are already familiar with various Storage solutions but for example want to know more about why you should consider using SAN technology – including benefits of SAN, Virtualisation and SAN or tips on Making Storage Highly Available, then **jump to page 8**

Storage

Storage

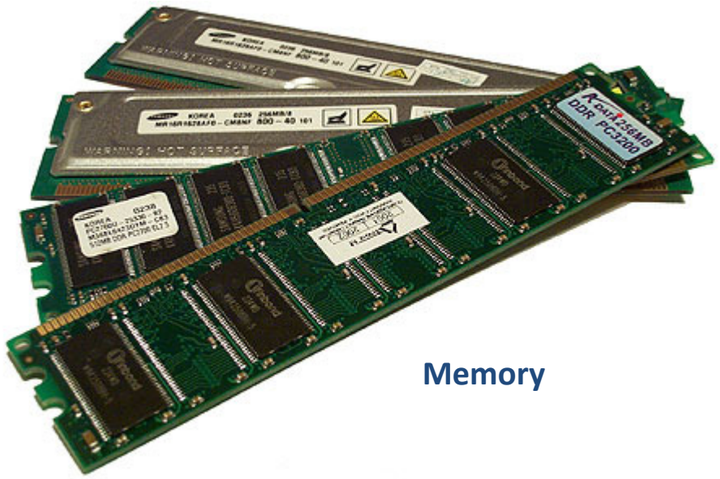

Memory

Definitions

Storage: The part of a computer that stores information for subsequent use or retrieval. Storage is frequently used to mean the devices and data connected to the computer through input/output operations - that is, hard disk and tape systems and other forms of storage that don't include computer memory and other in-computer storage.

The distinction between Memory and Storage is important here. Storage refers to storage devices and their media not directly accessible by the CPU, (secondary or tertiary storage) — typically hard disk drives, optical disc drives, and other devices slower than RAM but more permanent.

Memory: usually refers to semiconductor storage random-access memory (RAM), typically DRAM (Dynamic-RAM). Memory can refer to other forms of fast but temporary storage.

Historically, memory has been called main memory, real storage or internal memory while storage devices have been referred to as secondary storage, external memory or auxiliary/ peripheral storage.

The distinctions are fundamental to the architecture of computers. The distinctions also reflect an important and significant technical difference between memory and mass storage devices, which has been blurred by the historical usage of the term storage.

Storage

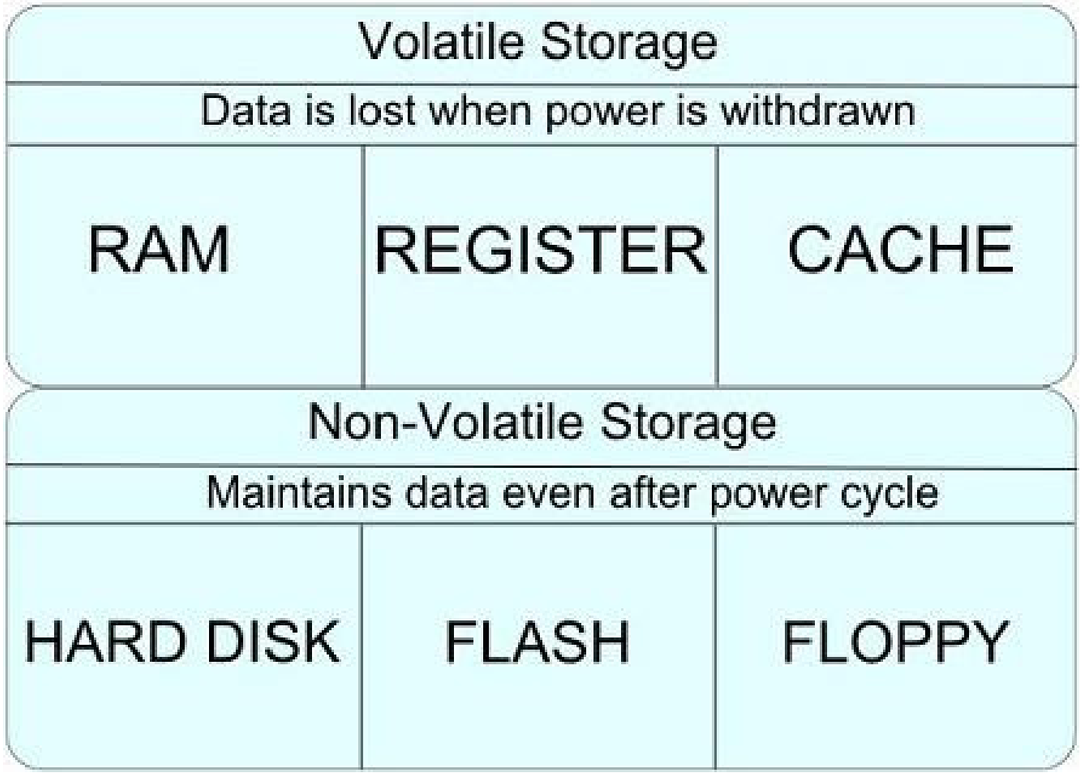

Primary and Secondary Storage

Primary storage: Primary storage (or main memory or internal memory), often referred to simply as memory, is the only one directly accessible to the CPU. The CPU continuously reads instructions stored there and executes them as required. Any data actively operated on is also stored there in uniform manner.

The most common primary storage is RAM. A key characteristic of RAM is that it is volatile, i.e. it loses the information when not powered.

Secondary storage: Secondary storage differs from primary storage in that it is not directly accessible by the CPU. The computer usually uses its input/output channels to access secondary storage and transfers the desired data to primary storage. Secondary storage does not lose the data when the device is powered down - it is non-volatile. Secondary storage is cheaper than primary storage.

In modern computers, hard disk drives are usually used as secondary storage. The time taken to access a given byte of information stored on a hard disk is typically a few thousandths of a second, or milliseconds. By contrast, the time taken to access a given byte of information stored in random access memory is measured in billionths of a second, or nanoseconds. This illustrates the significant access-time difference which distinguishes solid-state memory from rotating magnetic storage devices: hard disks are typically about a million times slower than memory.

Rotating optical storage devices, such as CD and DVD drives, have even longer access times. With disk drives, once the disk read/write head reaches the proper placement and the data of interest rotates under it, subsequent data on the track are very fast to access.

Some other examples of secondary storage technologies are: flash memory (e.g. USB flash drives or keys), floppy disks, magnetic tape and Zip drives. These operate even slower.

Storage

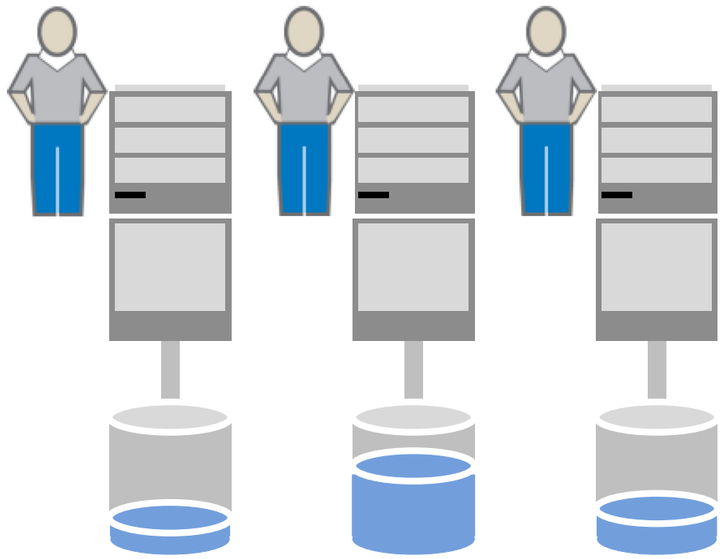

Direct Attached Storage – disks directly attached to each server

Secondary storage is connected to a computer typically in one of three ways:

1. Direct Attached Storage (DAS): is a traditional mass storage system that is either within or connected directly to a server, without a storage network in between.

Storage

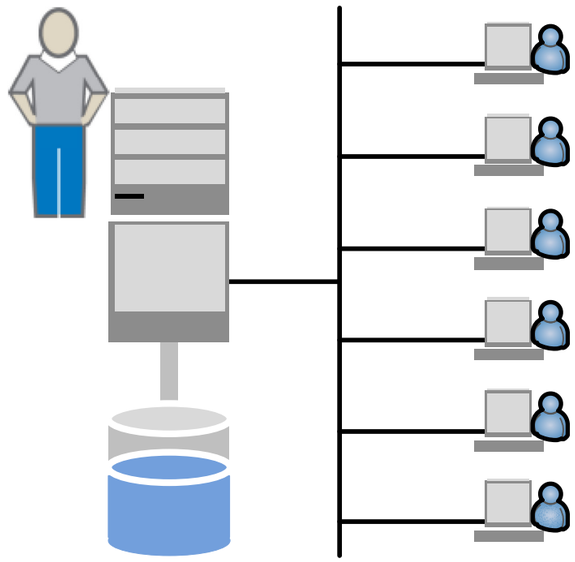

Network Attached Storage – mass storage accessed by PCs

2. Network Attached Storage (NAS): is mass storage attached to a computer which other computers (typically PCs) can access at file level over a network.

Storage

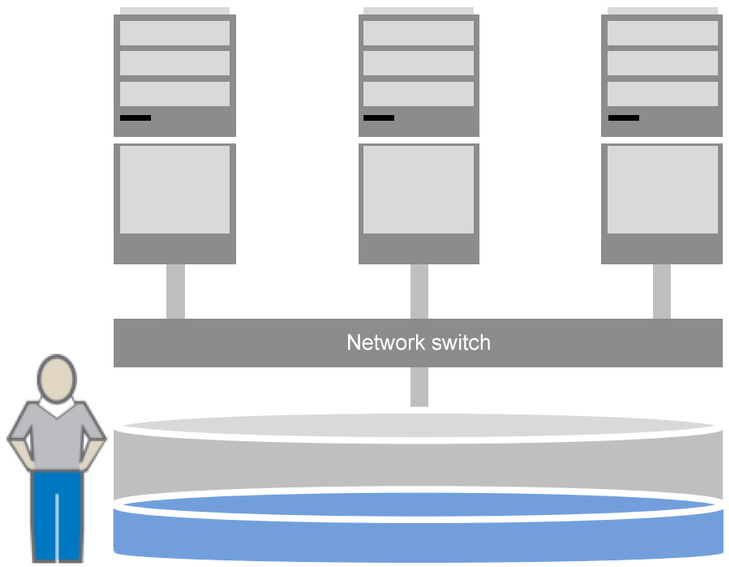

Storage Area Network – dedicated network that provides access to consolidated, block level data storage

3. Storage Area Network (SAN): is a dedicated network that provides access to consolidated, block level data storage. SANs are primarily used to make storage devices, such as disk arrays, tape libraries, and optical jukeboxes, accessible to servers so that the devices appear like locally attached devices to the operating system. A SAN typically has its own network of storage devices that are generally not accessible through the local area network by other devices. The cost and complexity of SANs has dropped in recent years to levels allowing wider adoption across both enterprise and small to medium sized business environments.

Storage

Problems with traditional storage – DAS & NAS and how a SAN addresses these problems

DAS

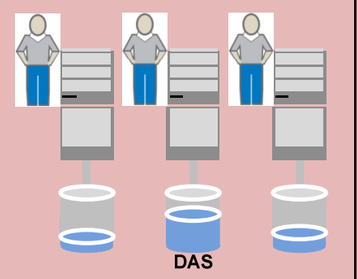

NAS

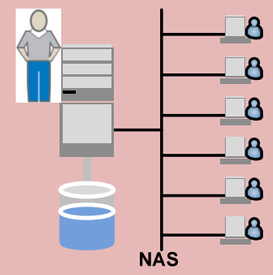

SAN

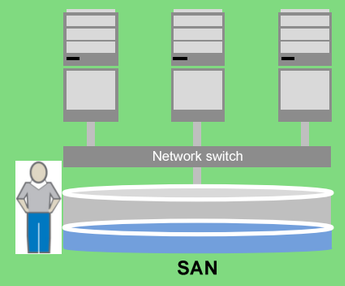

♦ Multiple islands of storage	⇒ Single storage island
♦ Uneven utilisation as data is unevenly spread across separate disk systems	⇒ Maximises utilisation
♦ Complex backup, often involving tape – multiple backup drives; long backup windows; media errors; etc.	⇒ Simple backup
♦ Complex to manage and inefficient – requires more Administration	⇒ Reduced complexity, less Admin
♦ Expensive continuity – any storage high availability solution needs to be repeated for each separate storage system	⇒ Affordable continuity
♦ Inefficient	⇒ Optimal efficiency
♦ Requires more physical space, power and cooling	⇒ Reduced space, power and cooling
♦ Requires commitment to the amount of physical storage at acquisition	⇒ ‘Pay as you Grow’ storage

Storage

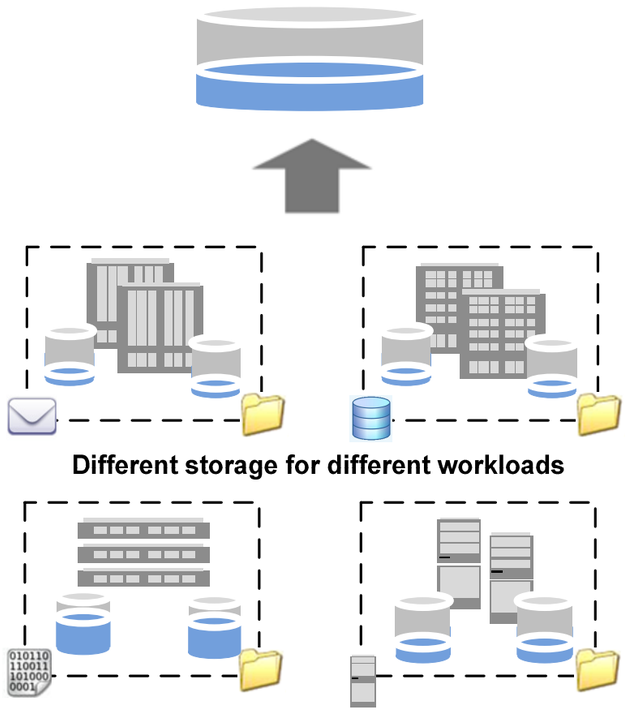

So what is a SAN and how does it solve these issues

A Storage Area Network (SAN) is a dedicated network that provides access to consolidated, block level data storage, and helps you:

Consolidate	⇒	Eliminate sprawl and maximise utilization
Manage	⇒	Reduce overhead
Protect	⇒	Reduce disruptions
Scale	⇒	Provision in minutes rather than days

Lower costs and increase productivity

Eliminate Sprawl and maximise Utilization

- ⇒ Handle all your storage needs with a single storage architecture – Eliminate individual DAS, NAS and file server silos
- ⇒ Optimize storage efficiency – maximise utilization
- ⇒ Regain control over budget – Reduce administration

Reduce Overhead

- ⇒ Consolidate and automate administration – Infrastructure, storage, application, and data
- ⇒ Speed time to deployment – Deploy in minutes, not hours or days
- ⇒ Integrate with existing Windows management – Spend time on productive tasks

Reduce disruptions; Provision in minutes rather than days

- ⇒ Near-instant backups of all data – Eliminate backup windows and application downtime
- ⇒ Application-aware backup to produce consistent backups – Exchange, SQL Servers, SharePoint and Virtualization
- ⇒ Simple restores – View previous backups and restore individual files and single mailboxes

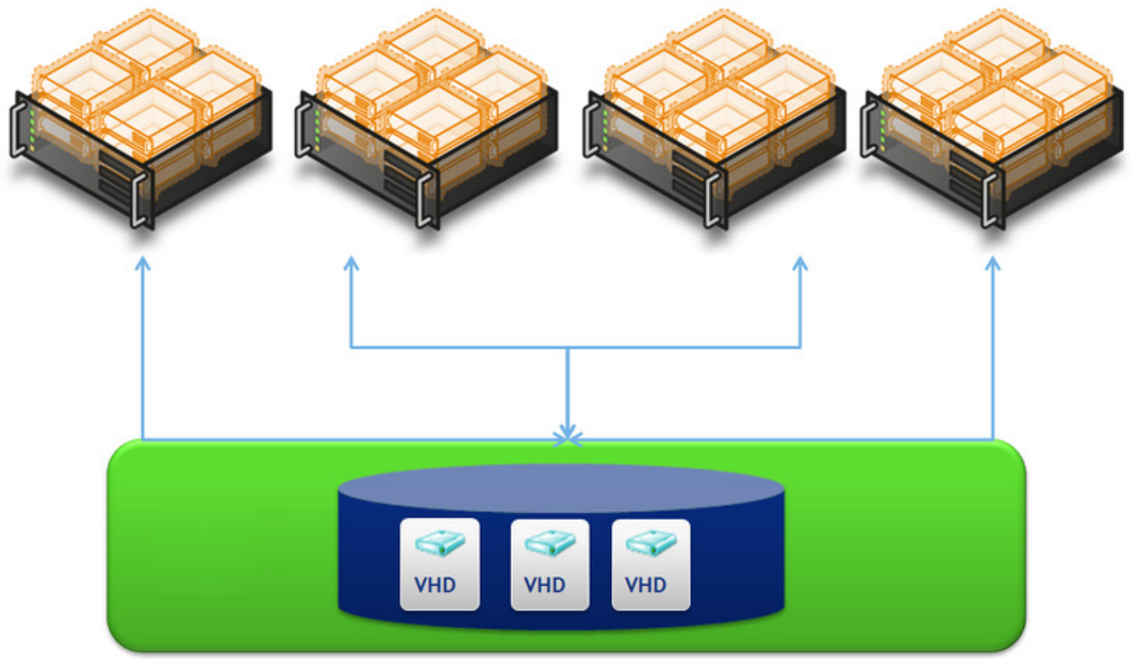

Four physical host computers running various virtual servers—all sharing a single storage array connected via a SAN

(VHD—Virtual Hard Disk)

Server Virtualisation – why use a SAN?

Most companies today have decided that if they are going to do a new server project, virtualisation is the way forward, as virtualization is one of the most effective methods to save costs, increase the availability, and improve agility within an organization's infrastructure.

Microsoft Hyper-V Server provides a simplified virtualization solution (other server virtualisation vendors include VMware and Citrix). Key features include dynamic memory, live migration, host clustering and Cluster Shared Volume (CSV) support – adding flexibility to consolidated servers.

Live Migration – allows you to move running virtual machines (VMs) from one Hyper-V physical host to another, without any disruption or perceived loss of service, providing support for dynamic load balancing, VM placement, high availability for virtualised workloads during physical maintenance, and reduced power consumption.

Cluster Shared Volumes (CSV) storage simplifies and enhances shared storage usage. It enables multiple Windows Servers to access shared SAN storage. CSV enables faster live migration and easier storage management for Hyper-V when used in a cluster configuration.

Although it is possible to implement server virtualisation using direct attached storage, to benefit from Live Migration – making workloads highly available – you require a storage area network (SAN). An alternative to Live Migration and for use with servers with direct attached storage, is a solution called **Double-Take Availability** – which effectively mirrors servers, apps and data between two servers, each with their own DAS. Should server A fail, server B kicks in.

Storage

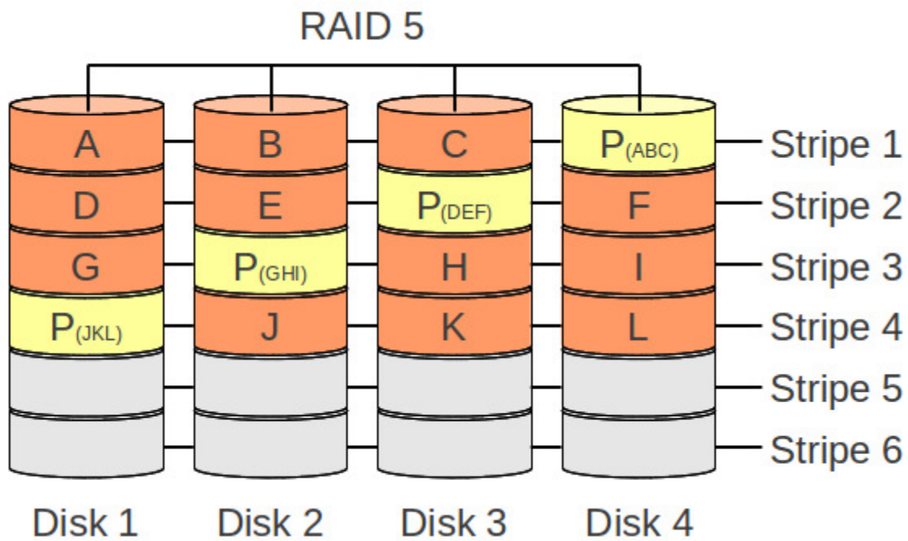

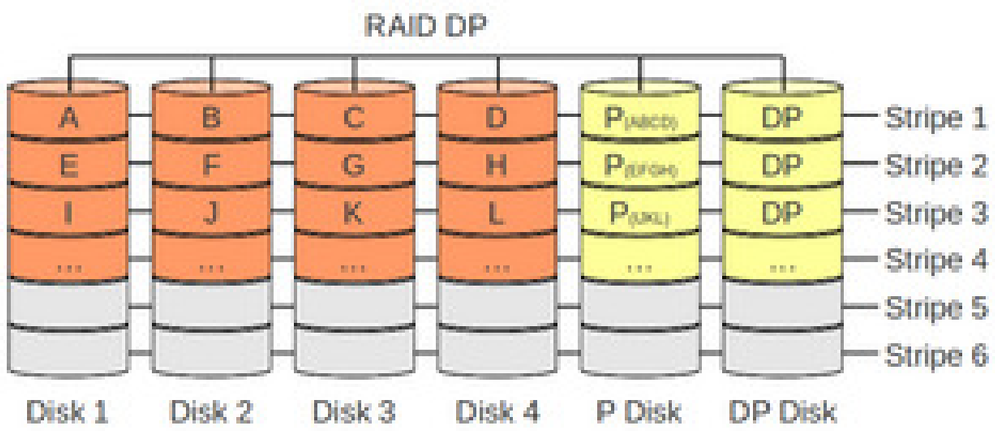

Making storage highly available

Clearly if you are going to put all your data into a single storage solution (as in the case of a SAN), you want to ensure it is highly available and doesn't suffer from single points of failure. Most storage arrays can be designed with redundant components:

Dual-controllers: The controller manages the physical disk drives and presents them as logical units. Dual-controllers can provide automatic failover in the event that one of the controllers fail, transparent to the server systems attached.

RAID: Is a technology that provides increased storage functions and reliability through redundancy. This is achieved by combining multiple disk drive components into a logical unit (Array), where data is distributed across the drives in one of several ways called "RAID levels". Many RAID levels employ an error protection scheme called "parity" which effectively provides fault tolerance as the array is not destroyed by a single drive failure.

RAID-DP (Double-Parity) or RAID 6: Provides fault tolerance of two drive failures; the array continues to operate with up to two failed drives. This makes larger RAID groups more practical, especially for high-availability systems. This becomes increasingly important as large-capacity drives lengthen the time needed to recover from the failure of a single drive. Single-parity RAID levels are as vulnerable to data loss as a RAID 0 array (without parity or mirroring) until the failed drive is replaced and its data rebuilt; the larger the drive, the longer the rebuild takes. Double parity gives time to rebuild the array without the data being at risk if a single additional drive fails before the rebuild is complete.

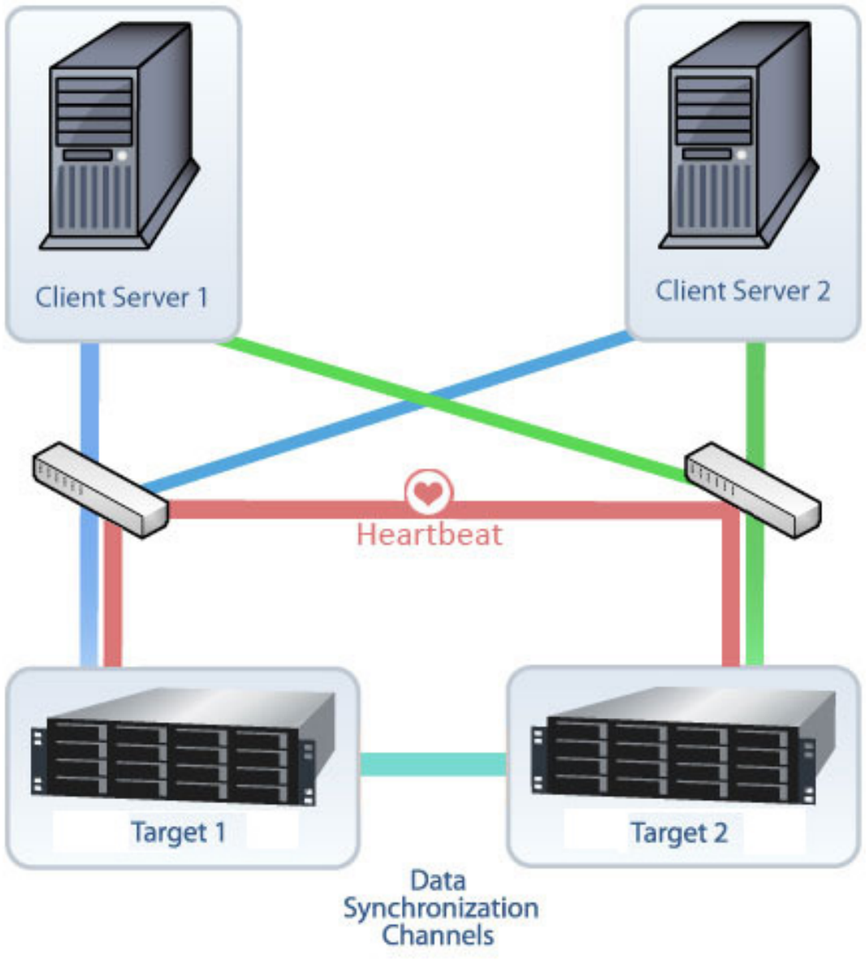

Making storage highly available, continued

Hot Spare: A drive physically installed in the array which is inactive until an active drive fails, when the system automatically replaces the failed drive with the spare, rebuilding the array with the spare drive included. This reduces the mean time to recovery (MTTR), but does not completely eliminate it. Subsequent additional failure(s) in the same RAID redundancy group before the array is fully rebuilt can result in data loss. Rebuilding can take several hours, especially on busy systems.

Dual SAN Switches: By using two or more SAN switches and connecting storage controllers, switches and physical servers using multiple patch leads providing multiple routes , you can ensure connectivity between servers and storage in the event that one of the SAN switches should fail.

Dual HBA Cards: The Host Bus Adapter (HBA) connects a host system (the server) to other network and storage devices. Implementing dual HBA cards in a server helps ensure connectivity to the SAN in the event a single HBA card should fail.

SAN-to-SAN replication: SANs tend to enable more effective disaster recovery processes. A SAN could span a distant location containing a secondary storage array. This enables storage replication either implemented by disk array controllers, by server software, or by specialized SAN devices. The purpose of replication is to prevent damage from failures or disasters that may occur in one location, or in case such events do occur, improve the ability to recover.

Storage

In conclusion

Clearly there are different options available when considering storage. Here we have tried to demonstrate some of the advantages of using a SAN over more 'traditional' methods – and given some guidance where high availability of data is important.

We have tried to convey some of the options – with their benefits and limitations – to assist you when considering which option is right for you.

Storage

Want to discuss your Storage Strategy or any other aspect of how technology can support your business?

How to contact Dave Graesser (Managing Director and senior consultant)

Office Phone	0370 350 2552
Email	dave@tekresponse.co.uk
LinkedIn	http://www.linkedin.com/in/davegraesser
Twitter	http://twitter.com/#!/DaveGraesser
Social Blog	http://davegraesser.blogspot.com/

Whether you have a Storage Strategy or the subject of Storage is new to you, we welcome your views and questions. As with all technology, it is important to review the options available and identify the potential implications and opportunities within your business. We cannot cover everything within this brief overview document – this is simply a snapshot of the areas to consider with storage technology.

As founder of Tek Response, an established multi-vendor accredited company I am proud of our history of delivering IT solutions which are individually tailored to our clients needs. A key goal in our business is to develop a long term partnership with clients and through our dialogue with them to understand their business needs and challenges and provide them with the advice and expertise they need when making choices about the appropriate technology for their business. I believe this is why today we work with such a diverse range of companies spanning many industry sectors.

My goal in producing these short guides is to help inform our clients a little about current technology trends; how these can improve productivity and performance, whilst reducing costs; and to perhaps help de-mystify some of the jargon and information overload which has become common in today's "communication rich" environment.

If you have any questions or would like to discuss your Storage Strategy or any other aspect of how technology can support your business, please do contact me. I hope this information has been useful and welcome your feedback into any improvements you would suggest or topics you would like to see covered.