



HIGH AVAILABILITY & STORAGE SOLUTIONS
ESI INFORMATION TECHNOLOGIES INC.

UNDERSTANDING NETWORK STORAGE

Martin Gauthier,
Associate Director, HA & Storage Solutions,
ESI Information Technologies Inc.

Montreal, Canada,
June 2003.

DEMYSTIFYING STORAGE

No one in the IT world can avoid it, at one point or another, storage becomes a serious issue. The exact issue varies from business to business. For some it may present itself as an ever increasing amount of data, which becomes impossible to backup within the available backup window. For others it may arise from an unmanageable spreading of data across multiple servers. Whatever the reason, everyone eventually faces an unavoidable explosion of their data.

The solution is to migrate to a networked storage environment. Unfortunately, the IT industry does not make things simple. Various vendors offer a multitude of products: DAS, SAN, NAS, and even iSCSI. But what are they?


First of all, one has to understand that these different products are not all of the same category, and therefore should not be compared to each other. Doing so is like comparing apples to oranges. To be more meaningful, one should compare DAS (direct-attached storage) and NAS (network-attached storage) to SAS (SAN-attached storage). As for SAN (storage area network), it should be compared to LAN (local area network) and WAN (wide area network). And what about iSCSI? It should be compared to another SCSI-embedding method, namely FC (fibre channel) rather than DAS, SAN and NAS.

DIRECT, NETWORK AND SAN-ATTACHED STORAGE


Contrary to NAS or SAS, DAS are storage devices that are connected directly to single or pair of servers, without the use of switches. The actual type of physical media used to connect devices, the protocol used to communicate (SCSI, FC, IP, etc...), and whether or not the disk controllers are located within the servers, do not matter. What makes DAS what it is, is the peer-to-peer architecture resulting from the absence of switches in-between it and the servers.



Otherwise, we are in the presence of NAS (networked-attached storage) and SAS (SAN-attached storage) devices when switches are used to connect storage devices to servers. As to differentiate between the NAS and SAS, one has to examine how data is accessed on the storage device. If data is accessed at the file-level, the storage device is a NAS. On the other hand, if data is accessed at the block-level, the storage device is a SAS.



Contrary to what is commonly thought, NAS and SAS devices cannot be differentiated by whether or not they are connected to IP switches or FC switches. Although it is true that NAS devices exclusively communicate using the IP protocol, SAS devices may also use that protocol to communicate. Such SAS devices are then referred to as iSCSI storage devices, where block-level SCSI commands and data blocks are embedded inside IP packets, as opposed to FC where SCSI commands and data blocks are embedded inside FC frames. The end result is the same: a server stores data on the SAS device, and communicate with that device as if it was physically located within the server.



iSCSI devices are often depicted as storage devices connected directly to the LAN. Although it is not the only way to connect iSCSI devices, especially now that most FC switches also support the IP protocol, it makes perfect sense to use the existing IP infrastructure. To be accurate, the portion of the LAN to which iSCSI devices are connected should be referred to as a SAN, or more specifically, an IP-SAN.



To really understand network storage, one has to understand the true nature of NAS devices, which unfortunately may be confusing to some people. NAS devices are made of two components: a server front-end and a storage back-end. The server front-end, or pair of servers (for redundancy), has two main functions: (1) to

provide network file-level access to block-level storage using various network sharing file systems (NFS, CIFS, etc...), and (2) to host the GUI (graphical user interface) utility through which the NAS device can be configured and managed. Hidden behind the server front-end is the storage back-end, which consists of either a DAS or a SAS device.




Some vendors offer the server front-end of NAS devices as a stand alone product. These servers front-end, referred to as NAS gateways, allows customers to add NAS functionalities to their existing SAS devices.

STORAGE, LOCAL AND WIDE AREA NETWORK

The most common mistake, when discussing about network storage, is to improperly use the SAN (storage area network) acronym. The latter is often incorrectly used to identify individual SAS (SAN-attached storage) devices rather than the actual storage network. In many ways, a SAN is comparable to a LAN (local area network) or a WAN (wide area network) in the sense that it represents the entire network dedicated to the interconnection of storage devices with servers.

Another common mistake is to automatically assume that having a SAN implies the use of FC (fibre channel) technology. The truth is that FC is just one of several protocols that can be used to communicate with SAS (SAN-attached storage) devices. FICON (fibre connect), for example, is another protocol commonly used for mainframe SAN.



IP (internet protocol), used for iSCSI storage devices, is also another protocol available for SAN. It is only recently, however, with the introduction of Gigabit Ethernet technology, that it became a realistic alternative to FC (fibre channel). Its greatest impact on the evolution of network storage will be in a few years when 10 Gigabit Ethernet technology becomes readily available.