IS MANAGEMENT 'IN YOUR FACE'

ABOUT YOUR INTERFACE STRATEGY?

Overland STORAGE™

TABLE OF CONTENTS

OVERVIEW

Companies working on interface strategies for the next two years are wondering which standards will win, which will die, or if they all will exist side by side. This creates pressure for IS/IT professionals to quickly develop a strategy that outlines to company executives the interfaces needed to move the organization forward and why.

Since early 2000, a "SAN" movement to provide SAN-like functionality over Ethernet infrastructure has taken hold. This has led to a plethora of standards activities resulting in the latest batch of alphabet soup standards (iSCSI, iFCP and FCIP). The current market is awash with different existing and proposed storage network connectivity options and will continue in this manner for the next 12 to 18 months as the market wrestles to find a clear winner, especially in the Internet Protocol (Ethernet)-based block level realm.

A tape automation interface strategy should outline plans for the next 24 months. Tape automation interfaces provide the ability to use a broad range of tape automation products that support both direct server attachments and shared storage attachments.

This paper will explore/discuss the following interfaces: ATA, SCSI, Fibre Channel, InfiniBand, Gigabit Ethernet and 10 Gigabit Ethernet.

INTERFACE APPLICATION DEFINITIONS

Within today's modern enterprise data center many different interfaces are used to meet specific applications requirements. As an example, today's disk array often contains two different interfaces. A host side (Front End) interface that is used to connect to a host or SAN / LAN, and a device interface (Back End) used to connect to the actual storage device – in this case, the disks.

DEVICE (subsystem) LEVEL INTERFACE DEFINITION

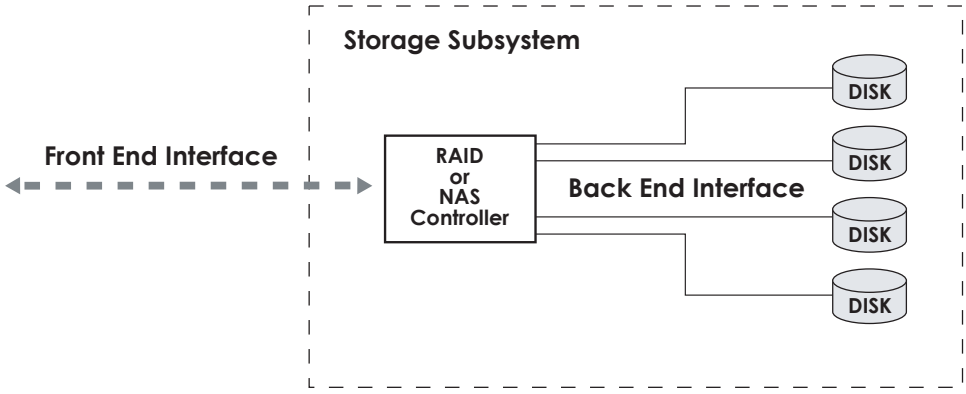

Front End interfaces such as 100/1000 BaseT Ethernet, SCSI or Fibre Channel provide storage-to-server interconnectivity within the data center. Connectivity between data centers within a campus (<10,000 meters) is sometimes done with Long Wave Fibre Channel over dedicated Fibre. Data centers that do not reside within a campus but are <80,000 meters apart may also take advantage of leased dark Fibre lines for a high price.

High bandwidth data center interconnection across the Metro Area Network (MAN) and the Wide Area Network (WAN) are possible, but at a substantial cost. Recent price quotes for dedicated bandwidth currently run about $1,000 per Megabyte per second per month excluding trenching and installation. Costs are expected to come down over time, but currently limit the use of high bandwidth pipes between data centers to a few select customers (banks, brokerage etc.) whose very business depends upon timely data replication and real time transactions.

END TO END DATA CENTER INTERCONNECTIVITY

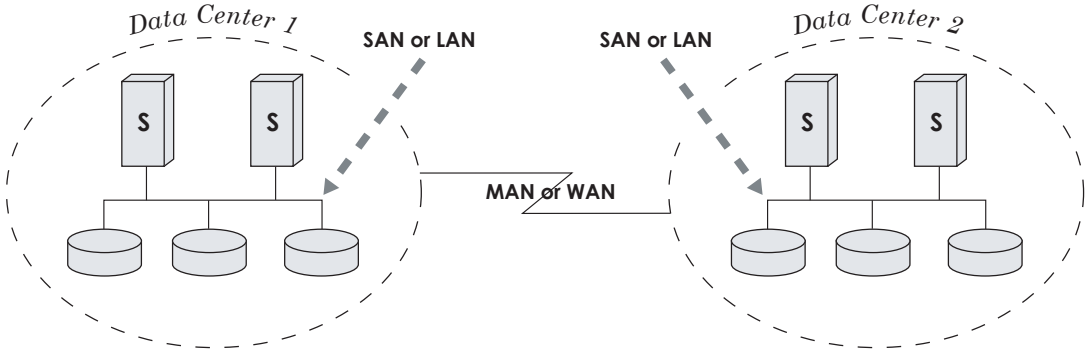

INTERFACE BRIEFS

The following section is intended to give a brief overview of the interface technology, along with its impact on the market.

ATA, SERIAL ATA

Advanced Technology Attachment or ATA will continue for the next three to five years as a disk level device interface for desktop, workstations and laptops. ATA is the evolutionary successor to the original IDE (Integrated Drive Electronics) hard disk interface popularized by the success of the IBM PC and the associated clone market in the early 1980s. Current ATA pricing is 75 to 80 percent cheaper than Fibre Channel or SCSI drives of equal capacity. Over the next few years, we will see more and more disk arrays incorporating ATA-based interfaces as a means to deliver more cost effective storage to a price sensitive market. If it becomes available at parallel ATA pricing, Serial ATA will simplify RAID controller mid-plane architecture. Serial interfaces simplify RAID controller architecture by reducing the board level routing nightmares associated with parallel buses. Serial buses also are easier and less expensive in switching environments.

SCSI (AS A SOFTWARE PROTOCOL)

Small Computer System Interface or SCSI has been adopted by all major OS vendors as the software protocol of choice for controlling block level access to data, and has become the mainstay of enterprise class storage systems. Fibre Channel, as an example, uses SCSI as an upper layer (FC-4 layer) protocol providing block level access to data. SCSI Command Descriptor Blocks (CDBs) will remain the key software protocol for the foreseeable future for enterprise class storage. All of the block level protocols for Gigabit Ethernet including iSCSI, iFCP and FCIP use SCSI CDBs as their command protocol.

PARALLEL SCSI

Parallel SCSI will continue to be the interface of choice for direct attach external storage. Today's dominant parallel SCSI interface is Low Voltage Differential (LVD). Current shipping product conforms to one of three specified data rates over 68 parallel wires. These include Ultra 80 (80 MB/second), Ultra 160 (160 MB/second), and Ultra 320 (320 MB/second). The maximum bus length of LVD SCSI is 12 meters.

High Voltage Differential (HVD) SCSI has fallen by the wayside, surrendering most of the enterprise class attachment and clustering market to Fibre Channel. There has been a substantial decrease in the order rate for HVD-based SCSI products outside of legacy applications due to lack of support for HVD SCSI beyond the 40 MB/second rate.

SERIAL SCSI

Serial SCSI as a low cost, high-speed serial interface is the next logical progression in SCSI's more than 20-year run. But the main issue with Serial SCSI is whether it can become a cost effective alternative to Parallel SCSI.

Serial SCSI will compete with Serial ATA at the low end for disk within the desktop or workstation. Serial SCSI will also compete with Serial ATA and Fibre Channel as the backend interface for disk arrays. Tape drive vendors may also embrace Serial SCSI, although few currently highlight this interface in their product roadmaps.

FIBRE CHANNEL

Fibre Channel will continue to erode SCSI library attachment as mid-range companies begin to share their backup resources through networking technologies such as Fibre Channel SAN.

However, Fibre Channel is at serious risk of losing the backend disk array bus market to ATA/IDE, Serial ATA and Serial SCSI. Cost is the driving factor as Fibre Channel disks command a 400 to 500 percent premium over lower cost disk interfaces. Current disk array vendors will be hard pressed to continue to demand these high prices as arrays offering equivalent speeds and feeds enter the market based on ATA and Serial ATA.

INTERFACE BRIEFS (cont.)

As Fibre Channel infrastructure companies bombard the market with low cost 2 Gb/second hardware, we will see Fibre Channel continue to be a dominant front-end interface over the next two to three years.

INFINIBAND

InfiniBand has significant potential to be the interface choice for blade level server clustering for several reasons. Low latency, a simplified serial interface, and well defined management provide blade server designers a fabric on which to build tightly coupled parallel processing system. However, before trying to grow to the next level of data center fabric, InfiniBand must gain acceptance as the de-facto interface for clustering.

The topology used for InfiniBand is a switched fabric form similar to Fibre Channel. The three key components of the InfiniBand fabric are:

- **HCA –** (InfiniBand) Host Channel adapter – used to connect a host system to the InfiniBand fabric.
- **Switch –** (InfiniBand) Switch – used to interconnect HCAs to HCAs or to TCAs, using information contained within the InfiniBand transmission to determine the destination port for the serial data stream.
- **TCA –** (InfiniBand) Target Channel Adapter used to connect non-initiating devices to the InfiniBand fabric

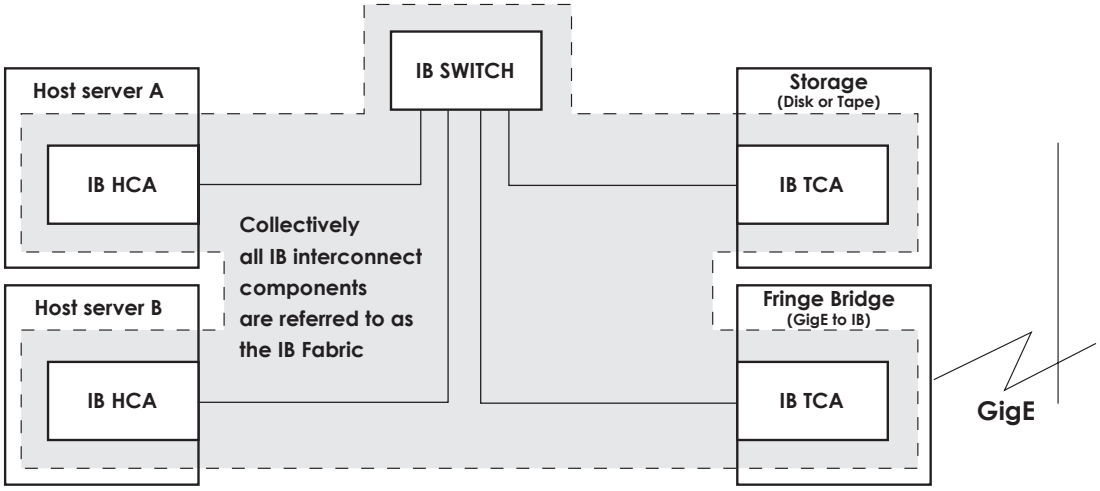

HCA technology is beginning to appear in the form of a PCI or PCIX daughter card that is inserted into a host server's existing I/O bus, similar to how one would install a Fibre Channel Host Bus Adapter or GigE NIC card.

Native (or built on the motherboard), HCAs are not likely to appear until mid- to late 2003.

The first real use for InfiniBand will be by the server vendors to perform high performance clustering of mul-

tiple servers running large database applications. The current InfiniBand specification defines three implementations of InfiniBand:

1X = 500 MB/second full duplex – (1 stand alone InfiniBand channel)

4X = 2 GB/second full duplex (4 InfiniBand channels working in parallel)

12X= 6 GB/second full duplex (12 InfiniBand channels working in parallel)

InfiniBand's low latency feature, Remote Direct Memory Access (RDMA), is the key value for server vendors. The advantage of RDMA is in its ability to tightly couple memory between physically separated processors, allowing for high performance parallel processing to occur by mitigating the impact of data movement between processor cards.

The following is a brief timeline as to where the IB industry is and will be over the next few years.

Currently HCAs, TCAs and Switches are available from several different companies, and rudimentary interoperability testing has begun.

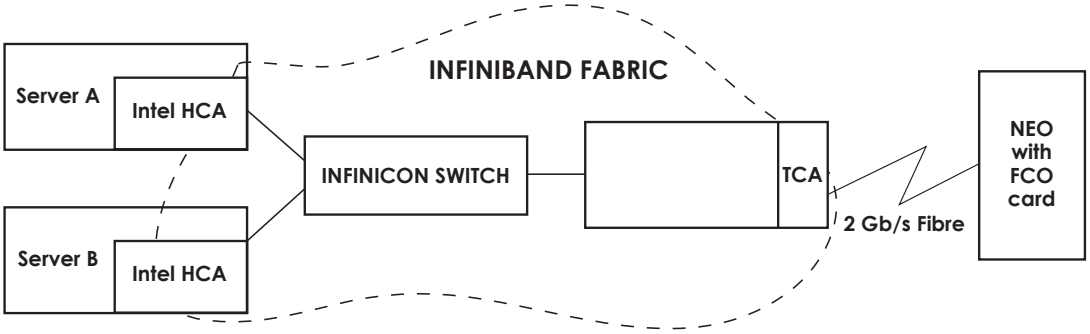

In summary, the InfiniBand "fabric" is here to stay. It will initially be deployed as a means for high-speed server clustering running single image database software (Oracle 9, DB2) on multi-server configurations. Long term, InfiniBand will be native as the edge connector for blade level servers.

If the cost of InfiniBand does not come down, it will face a rough uphill road to become a mainstream interface in the storage industry as this space is currently well served by Fibre Channel and Gigabit Ethernet.

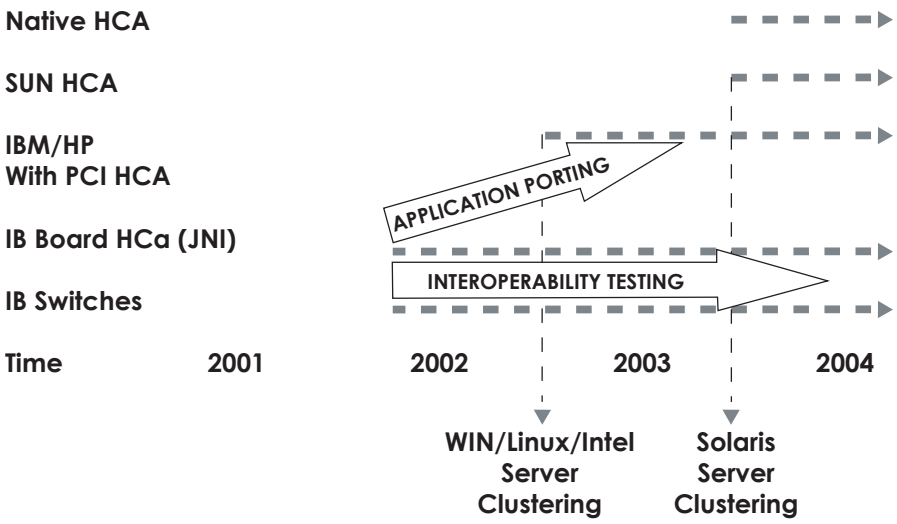

INTERFACE BRIEFS (cont.)

GIGABIT ETHERNET (GigE)

Gigabit Ethernet (GigE) will have a profound effect on storage going forward. GigE provides the hardware and TCP/IP foundation for several higher level protocols associated with data movement and backup such as the current Network Data Management Protocol (NDMP) protocol, as well as the emerging iSCSI, FCIP, and iFCP protocols.

Three different styles of LAN-based SANs are emerging as dominant. The first is a dedicated LAN whereby an existing backup server "pulls" the data from other servers and NAS devices using conventional file level protocols such as NFS and CIFS over Gigabit Ethernet. The back-up server then pushes the data out over direct attached SCSI to tape libraries for backup and archive purposes.

The second style of SAN functionality over a LAN is through the use of the NDMP protocol.

The third style of backup LAN will be iSCSI. This will be deployed as a dedicated storage LAN having SAN-like block level data movement capabilities over copper GigE between iSCSI RAID arrays, iSCSI enabled host and iSCSI attached libraries.

The following is a brief overview of a GigE system and its components. All of the current proposed storage over IP standards (iSCSI, FCIP, iFCP, NDMP) share a common hardware platform in addition to TCP/IP as a low-level protocol. The following drawing illustrates the key building blocks and their relationship.

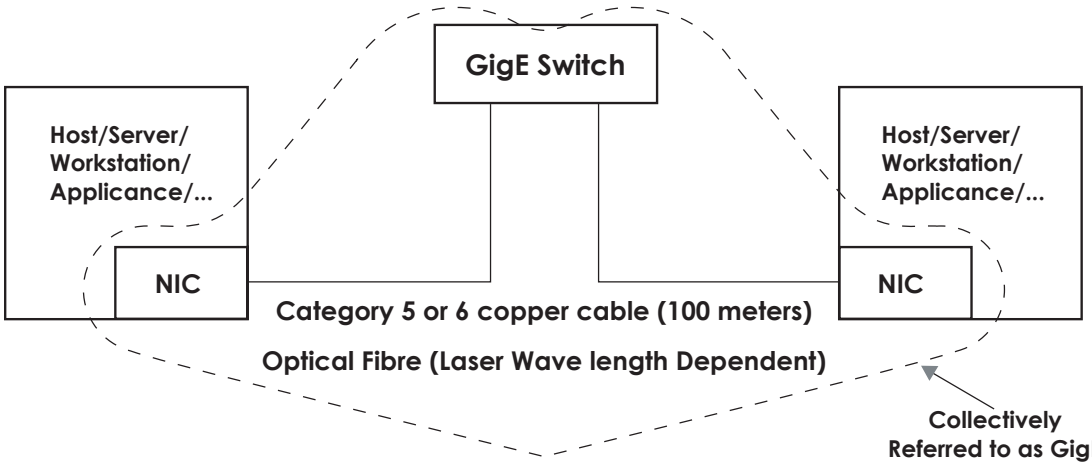

NIC – Network Interface Card – Provides the hardware bridge between the host's native bus and the 100/1000 Base-T Ethernet serial interface. Current low end (<$300) GigE NICs rely heavily on the host system's CPU to perform TCP/IP stack processing. This limits throughput to <40 MB/second and fully consumes a 1 GHz Pentium 3. Second generation GigE NICs will take advantage of TOE (TCP Off-load Engine) technology and improve sustainable throughput in to the 80-90 MB/second range.

Switch – Provides routing between ports on the switch based on IP addresses contained within the TCP IP packet.

Cabling – GigE provides for several different media types depending on the distance one wants to run. Simple CAT 5 E & CAT 6 copper cabling will provide up to 100 meters (or "run") from port to port. Short wave lasers takes it to 500 meters, while various long-wave lasers take it to over 80,000 meters.

Hosts / Servers / Appliances – For any of these systems or subsystems to take advantage of the Ethernet hardware and TCP/IP transport layer, they require a higher level protocol layer to define intra packet formats and bit usage. Hence, the need for iSCSI, iFCP, FCIP and NDMP.

Note: For devices (Hosts & Targets) to communicate with each other, they must speak the same language for the communication to be efficient. In other words, iSCSI hosts talk to iSCSI targets, while NDMP servers talk to other NDMP servers over GigE etc.

One key advantage that 1 GigE has over previous generations of Ethernet (10 and 100Base-T) is that it permits only switched traffic flow. Until just recently, most 10/100 LANs being deployed were of the hubbed variety, resulting in packet collision. That was followed by re-transmission during times of heavy LAN traffic, making them less than ideal from a storage perspective. GigE does not permit hubs and has eliminated delays caused by collision, yielding a more predictable throughput than its hubbed predecessors.

GigE will see deployment to the desktop starting in late 2002 for those applications such as document scanning that require tremendous amounts of data be moved between clients and file servers or NAS boxes.

NETWORK DATA MANAGEMENT PROTOCOL (NDMP)

NDMP is an established market for backup of high-end NAS filers. Originated by Network Appliance as a means to prevent the need for third party software, NDMP is a higher level protocol for Ethernet.

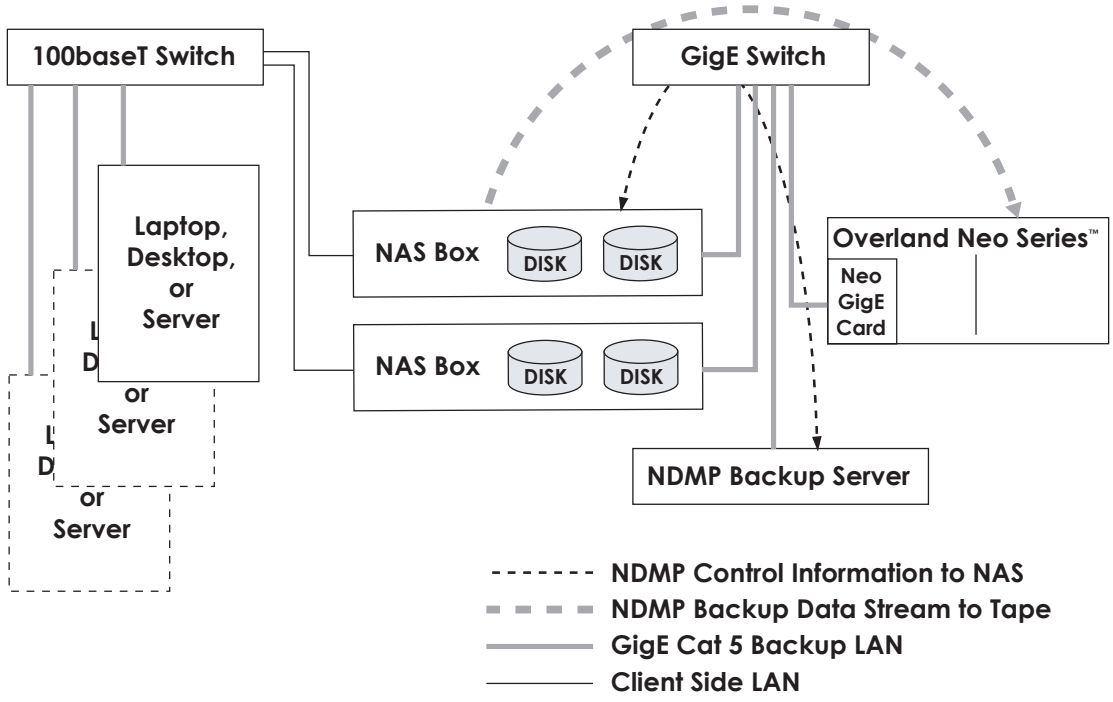

INTERFACE BRIEFS (cont.)

A tape library that supports NDMP version 4 will enable data to be transferred directly from the NDMP-compliant NAS device to the library under the control of the backup server running NDMP client software.

INTERNET SCSI (iSCSI)

iSCSI is a critical emerging protocol for Ethernet that will provide a means for block level delivery over a historically file-orientated interface. The iSCSI protocol will have tremendous impact on storage over the next three to five years as the Internet Engineering Task Force is expected to vote on the draft standard at the end of 2002.

IBM has supplied iSCSI RAID arrays for almost two years. Other major RAID vendors also are expected to introduce iSCSI attached RAID arrays by year's end.

iSCSI host attachment will come from both the conventional Fibre Channel / SCSI HBA camp, as well as the NIC camp.

FCIP FIBRE CHANNEL OVER IP (FCIP)

Fibre Channel (encapsulated within) IP (FCIP) provides connectivity between islands of Fibre Channel SAN across the MAN WAN using TCP/IP.

FCIP is a TCP/IP-based tunneling protocol for connecting geographically distributed Fibre Channel SANs to each other over IP. FCIP relies upon standard GigE hardware for switch functionality and routing through the IP LAN/MAN/WAN. FCIP bridges encapsulate incoming Fibre Channel streams from Fibre Channel switches and provide for connectivity within the data center. FCIP devices should be thought of as simple protocol bridges between Fibre Channel and GigE, providing only the necessary hardware and firmware to encapsulate Fibre Channel into TCP/IP packets and vice versa. MAN / WAN link speeds can be low and latencies can be high requiring timeout adjustments to be made in both SAN islands to allow for operation.

FCIP EXAMPLE

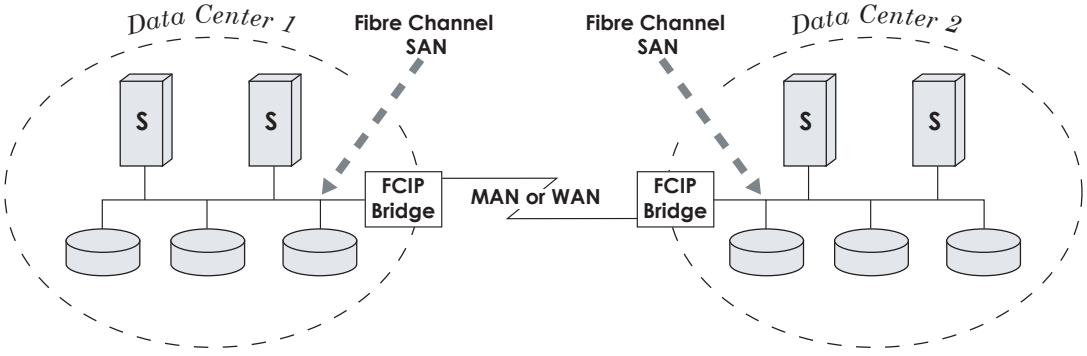

INTERNET FIBRE CHANNEL PROTOCOL (iFCP)

Internet Fibre Channel Protocol or iFCP provides SAN-like fabric functionality based on Ethernet switching. Replacing Fibre Channel switches with GigE hardware provides a lower cost "fabric".

iFCP is a TCP/IP based protocol allowing a "SAN" to be created with Fibre Channel end devices using all GigE switching and routing in between. iFCP devices provide both switching and bridging functionality. iFCP devices are also referred to as "Gateways" devices.

iFCP's main competition will be from the iSCSI market. The main vendor for iFCP, Nishan Systems, has addressed this issue by incorporating iSCSI ports on their iFCP switch. The following is an iFCP IP Fabric example.

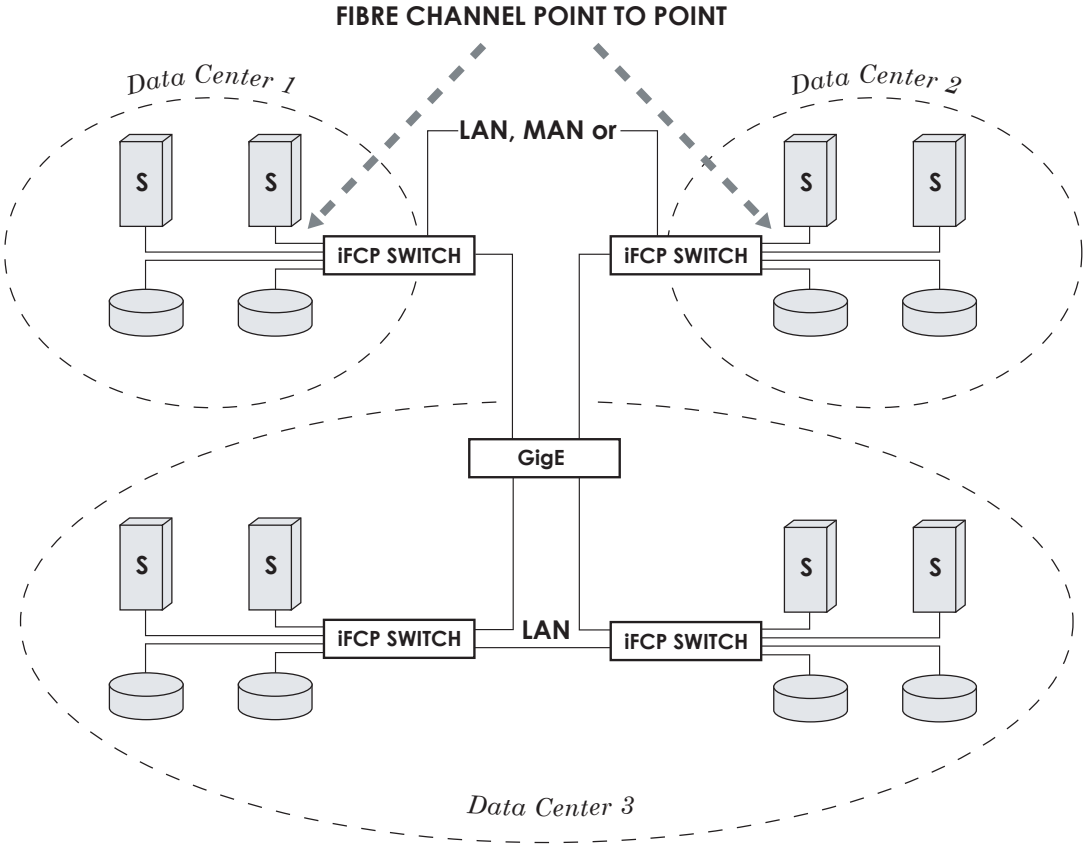

INTERFACE BRIEFS (cont.)

10 GIGABIT ETHERNET (10 GIGE)

In June 2002, the IEEE standards body ratified 802.3ae, the 10 Gigabit Ethernet (10 GigE) standard. 10 GigE will become the backbone (or main pipe) of choice for large data centers by mid 2004, and smaller data centers will adopt this protocol as costs to implement it decrease.

10GigE also will become the point of complete hardware convergence between block level SANs and file level NAS as 10 GigE pipes will concurrently support iSCSI, iFCP, NDMP, as well as the file level-like NFS and CFS. Products from Cisco and other high-end networking vendors are trickling into the markets that use 10 GigE as a switch to switch pipe. The cost for a single 10 GigE port is more than $10,000.

As costs for 10 GigE ports drop and 10 GigE switches become available, data centers will begin to deploy 10GigE fabric capable of 1 GB/second as their "corporate" backbone. Over time, GigE will move to the desktop providing 100 MB/second bandwidth to clients on the corperate LAN. Last mile costs should decrease then, providing 100 Mb/second access across the MAN / WAN.

EARLY 10 GigE DEPLOYMENT

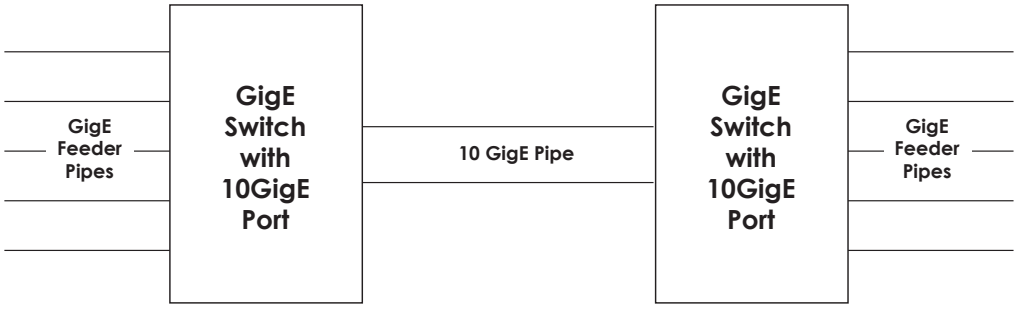

MID 2004 GigE LARGE DATA CENTER DEPLOYMENT

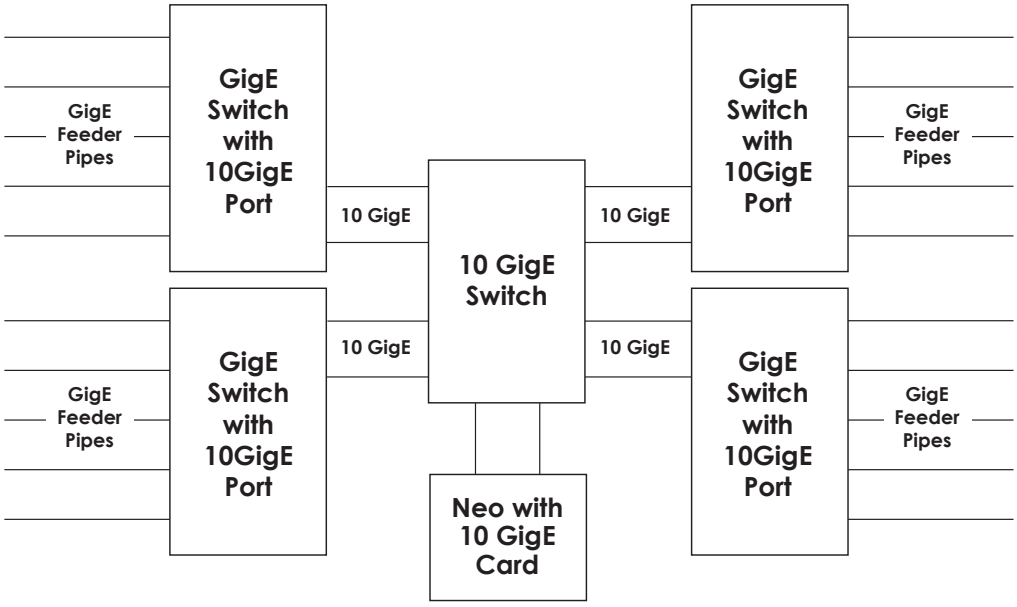

CONCLUSION

Overland Storage is a leading global supplier of innovative hardware and software storage solutions for mid-range computer networks. Creating tape automation interface strategy of its own that makes the best use of the most appropriate standards has been a high priority for the company.

The tape automation interface market is evolving responding to the dynamic needs of storage. Tried and true parallel LVD SCSI will continue to be the interface of choice for the next two years for the low-end stand-alone drive and loader markets.

Low cost GigE infrastructure (Switches, NICs, Routers, Cat 5 & 6 cabling) will enable small- to medium-sized companies to create dedicated backup SANs based on GigE LAN hardware for use within the data center. This will fuel both the existing NAS and emerging iSCSI markets.

The GigE NAS market running existing file level protocols such as CIFS and NFS and NDMP for backup will continue to be a robust market for midrange companies seeking simple solutions to the rapidly growing file based data.

Overland Storage believes that a new market based on GigE running iSCSI will emerge, providing lower performance (75-85 MB/second) SAN-like functionality at 30 to 40 percent of the cost of Fibre channel for small- to mid-range companies.

In addition, 1- and 2- Gb/second Fibre Channel will be the data center fabric of choice for the next two to three years. 10GigE will provide the pipe convergence of storage. Block level data (iSCSI, FCIP and iFCP), as well as file level data, will flow concurrently over the 1-GB/second 10GigE corporate backbone. Overland Storage will be there with products that seamlessly integrate with networked storage, providing the highest levels of data protection.

WORLDWIDE HEADQUARTERS
4820 Overland Avenue
San Diego, CA 92123 USA
TEL 1-800-729-8725
1-858-571-5555
FAX 1-858-571-3664
EMAIL sales@overlandstorage.com

UNITED KINGDOM (EMEA OFFICE)
Overland House
Ashville Way
Wokingham, Berkshire
RG41 2PL England
TEL +44 (0) 118-9898000
FAX +44 (0) 118-9891897
EMAIL europe@overlandstorage.com

FRANCE OFFICE
Overland Storage
126 rue Gallieni
92643 Boulogne Cedex France
TEL +33 (0) 1 55 19 23 93
FAX +33 (0) 1 55 19 25 02
EMAIL europe@overlandstorage.com

GERMANY OFFICE
Humboldtstr. 12
85609 Dornach Germany
TEL +49-89-94490-214
FAX +49-89-94490-414
EMAIL europe@overlandstorage.com

ASIA PACIFIC REP. OFFICE
30 Robertson Quay, #02-10
Singapore, 238251
TEL 65-6839-3510
FAX 65-6738-3008
EMAIL asia@overlandstorage.com

WWW.OVERLANDSTORAGE.COM