TECHNOLOGY BRIEF

March 1999

Compaq Computer Corporation

ISSD Technology Communications

CONTENTS

Compaq E2000 Platform Architecture: Industry-Standard Enterprise Computing Technologies

The E2000 Platform Architecture represents the efforts of Compaq and its industry partners to bring the flexibility, scalability, and low cost of standards-based platforms to all parts of enterprise computing. With the Compaq E2000 Platform Architecture, customers no longer have to choose between power and flexibility in their enterprise computing environments. High availability, fault tolerance, and reliability, which were once the domain of proprietary mainframe computers, can now be obtained by distributed networks of standards-based servers, storage, and communication interconnects. Standards-based solutions reduce the costs of acquisition and support and expand customers' choices in providers of systems, applications, and services. The standards-based Compaq E2000 Platform Architecture also provides a common platform to foster industry cooperation and innovation.

Please direct comments regarding this communication to the ISSD Technology Communications Group at this Internet address: TechCom@compaq.com

COMPAQ

NOTICE

Compaq E2000 Platform Architecture: Industry-Standard Enterprise Computing Technologies
First Edition (March 1999)
Document Number ECG063/0399

INTRODUCTION

In an effort to reduce computing costs and improve competitiveness, enterprise customers are shifting from centralized, monolithic mainframe systems to distributed networks of servers, storage, and communications systems.

The Compaq E2000 Platform Architecture is a distributed computing architecture based on industry-standard components, including symmetric multiprocessing (SMP) servers, high-performance networking technologies, and storage systems, connected by scalable system area network interconnects. These robust systems are tightly integrated with industry-standard operating systems and enterprise-caliber applications capable of functioning in distributed computing environments.

This technology brief describes the standards-based technologies in the Compaq E2000 Platform Architecture. These technologies are organized according to the needs of enterprise customers: high scalability, high availability, and manageability.

HIGH-SCALABILITY TECHNOLOGIES

The E2000 Platform Architecture allows scalability in all aspects of enterprise computing systems, including processing, memory, and input/output (I/O), including storage and communications. The architecture allows one or more industry-standard SMP servers to be connected with robust network storage using high-speed, reliable communication interconnects. The technologies that enable this highly scalable architecture include

- SMP servers: Compaq ProLiant servers and Compaq AlphaServers
- Industry-standard clustering
- Virtual Interface Architecture
- ServerNet
- Gigabit Ethernet
- Fibre Channel-attached storage

Symmetric Multiprocessing Servers

The demands of distributed computing greatly expand the role of PC servers in enterprise computing environments. Industry-standard SMP servers are assuming enterprise computing responsibilities, such as application reliability and fault tolerance that were once administered only by proprietary mainframe systems.

SMP servers divide system resources and tasks among multiple processors that share memory and run on the same copy of the operating system. The major advantages of SMP servers are (1) low communication overhead between the processors, and (2) the ability to add processors without substantially increasing the total cost of ownership.

Compaq ProLiant servers and AlphaServers form the SMP building blocks of the E2000 Platform Architecture. These servers contain powerful processors based on the Intel Architecture and Digital Alpha processor platforms, respectively.

Compaq ProLiant SMP Servers

Compaq ProLiant SMP servers are based on the Intel Architecture 32-bit (IA-32) technologies: the Pentium® II Xeon™ processor and most recently, the Pentium® III Xeon™ processor. The Compaq ProLiant Pentium II Xeon servers provide business-critical performance that surpasses comparable 4-way and 8-way Pentium Pro-based systems. The features of the Compaq ProLiant Pentium II Xeon servers are shown in Figure 1.

Compaq IA-32 Symmetric Multiprocessing Servers

Expansion optimized for the ultimate flexibility	 **ProLiant 7000** • Up to four Pentium II Xeon Processors • Smart 3100ES Controller • Push Button PCI Hot Plug • 8 GB memory support • Internal storage to 218.4 GB (SCSI) • External storage to 1.65 TB (SCSI) • Hot-plug redundant fans • Hot-plug redundant power supplies (3 max) • Redundant processor power modules	 **ProLiant 6000** • Up to four Pentium II Xeon Processors Smart 3100ES Controller option • 1 ISA Slot, four 32-bit/33-MHz PCI slots, five 64-bit/33-MHz PCI Slots • Up to 8 GB ECC EDO RAM • Internal storage to 218.4 GB (SCSI) • External storage to 1.8 TB (SCSI) • Standard redundant processor power modules • 14U when rack-mounted
Density optimized for multi-server environments	 **ProLiant 6500** • Up to four Pentium II Xeon Processors • Up to 8 GB of ECC EDO RAM • Optimized for Racks - 7U, footprint of 2.5 sq.ft. with rack-modular design • Five 64-bit/33-MHz PCI Hot Plug slots and one shared PCI/ISA slot • Hot-plug redundant fans • Hot-plug redundant power supplies (2 max) • Redundant processor power modules	 **ProLiant 5500** • Up to four Pentium II Xeon Processors • Smart 3100ES Controller option • Up to 4 GB ECC EDO memory • Six PCI 64-bit/33-MHz slots and one shared PCI/ISA slot • Internal storage to 127.4 GB (SCSI) • External storage to 1.87 TB (SCSI) • Hot-plug power supply • Optional redundant fans

Figure 1. Intel Architecture 32-bit Compaq SMP servers for mission-critical enterprise applications.

IA-64 Processor-Based Servers

Intel Architecture 64-bit (IA-64) processor technology promises to add a new level of performance and additional scalability to Compaq's existing line of industry-standard servers. Compaq's industry-leading systems based on IA-32 technology provide a straightforward migration path to IA-64. Intel expects to deliver the first IA-64 processor, code-named Merced, in mid 2000.[1] Compaq will be the leading supplier of IA-64 servers.

[1] See white paper *Compaq Servers: Enterprise Class Systems Leading the Way to Deschutes and Merced*, document number ECG078/0498, for more information.

Superscalar – processors that use several execution units to execute more than one instruction at a time.

RISC – reduced instruction set computing

Alpha Processor-Based SMP Servers

Compaq AlphaServers are based on the Alpha processor, which is a high-performance, 64-bit, *superscalar, RISC* processor originally developed by Digital Equipment Corporation. Since its introduction in 1992, the Alpha processor has been the industry's fastest processor. The latest generation Alpha processor (EV6 or 21264) was introduced in 1998. The most significant features of the 21264 processor are its spectacular memory addressing capacity, its ability to perform out-of-order execution at a very high rate, and its performance of floating point tasks. Due to its success, Alpha is the only 64-bit processor supported by Microsoft Windows NT, Digital OpenVMS, and Digital UNIX.

The Alpha processor roadmap is shown in Figure 2. The Alpha 21264 has 15.2 million transistors (6 million in the CPU core) on a .35-micron, six-layer metal CMOS-6. At the end of 1998, its clock speed was 575 MHz with measured peak performance levels of 30 SPECint95 and 50 SPECfp95. Next, the Alpha 21264 processor will be re-implemented in a 0.25-micron process and then in a 0.18-micron process. The 0.25-micron version will have a clock speed of over 750 MHz, and the 0.18-micron version will have a clock speed of over 1 GHz. The next two major generations of Alpha are the EV7 (21364) and the EV8 (21464) processor. The EV7 will be implemented in a 0.18-micron, six-layer metal process; this integrated design will have 100 million transistors and will start with a clock speed of over 1 GHz. The EV8 will be implemented in a 0.13 micron process and will have a clock speed over 1.4 GHz.

Alpha product designations:
EV – Extended VAX

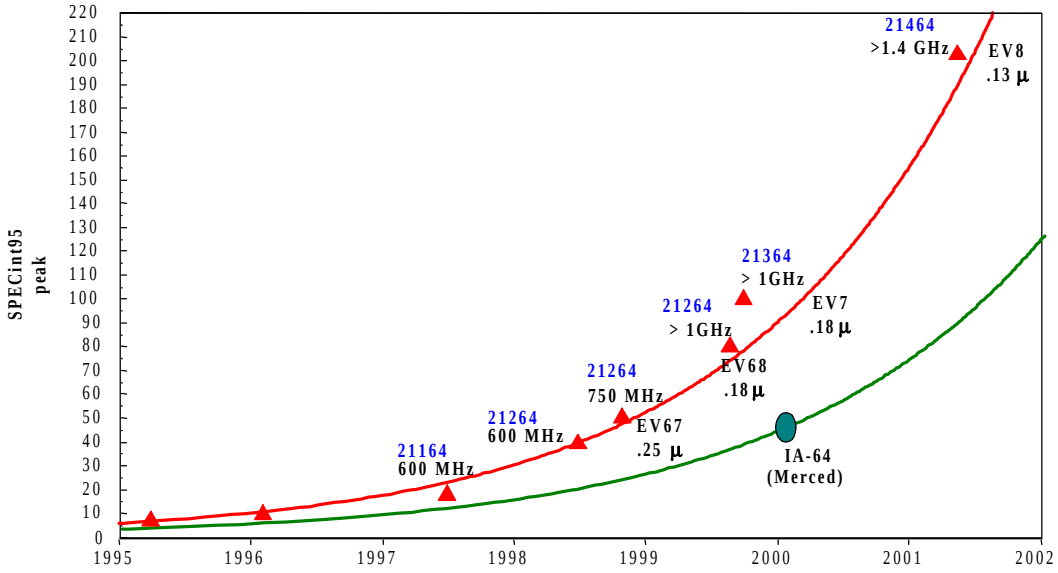

Figure 2. Alpha processor roadmap.

Compaq remains committed to the 64-bit AlphaServer line because:

- The Alpha processor will maintain its leadership and remain the clear choice for 64-bit computing.
- Compaq's experience with 64-bit processing will allow it to take advantage of the massive 64-bit application development that will be spurred by the introduction of the first IA-64 processor (Merced).

Compaq's AlphaServer GS series, based on the Alpha 21264 processor, continues to demonstrate superior performance and scalability in enterprise markets such as data warehousing and business intelligence. The first members of the Alpha GS series, the AlphaServer GS60 and GS140, provide up to 2.5 times the application performance of the AlphaServer 8200 and 8400 systems, which are based on previous-generation Alpha processors. The Compaq AlphaServer GS60 and GS140 servers feature

- Up to 14 Alpha processors, each with its own data and address paths
- 28 GB of error-correctable memory
- 144 individual PCI slots to configure up to 85 TB of Ultra SCSI storage (with future support for Fibre Channel storage)
- System and I/O bandwidth normally associated with more costly mainframe systems
- Excellent application performance and price/performance for users who need to solve large-scale computing problems
- An exceptional availability profile, with both UNIX and OpenVMS standard RISC platforms
- Flexibility to grow the computing environment, while managing it as a single entity and maintaining operations and management simplicity

Industry-Standard Clustering Technology

Cluster – a group of loosely coupled servers (each capable of operating independently) that work collectively as a single system.

In the past, only very expensive, proprietary mainframe systems could deliver the levels of performance, reliability, and scalability required for enterprise computing. Today, clustering technology allows customers to connect two or more SMP servers, called nodes, together to apply increasing numbers of processors to environments that require extensive computing resources.

Currently, there are two basic types of clusters: *application failover* and *parallel application*. In the 2-node application failover cluster (Figure 3 left), an application running on one server switches to a second server if the first server fails. In the parallel application cluster (Figure 3 right), all servers simultaneously run a copy of each application. Client requests are divided among the servers, which exchange information about which instance of the application will work on which data for each user. If one server fails, the jobs it was processing are distributed across the remaining servers.

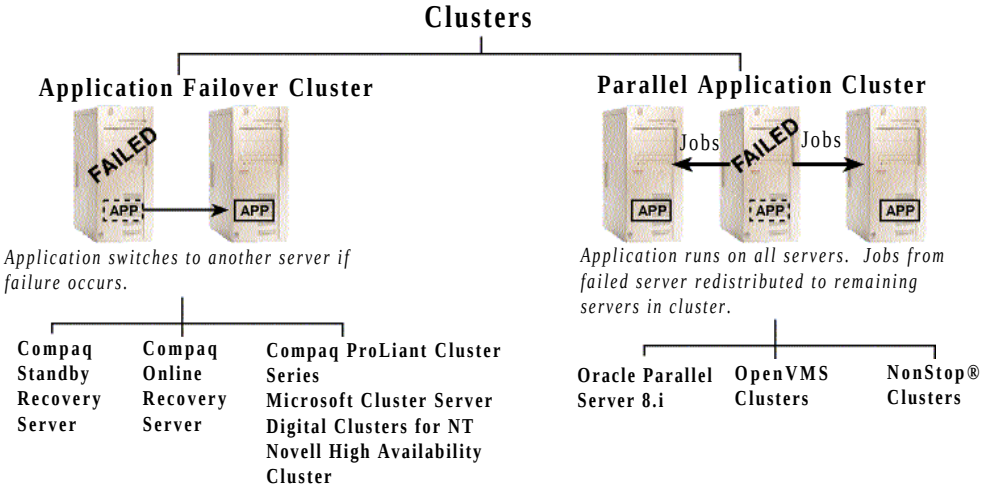

Figure 3. Cluster types

The benefits of parallel application clusters are substantial:

- Processing capacity and I/O bandwidth can be scaled by adding relatively inexpensive servers to the cluster.
- Clustered servers "back each other up" in case of failure, increasing availability.
- Clusters offer manageability by providing a single system image to the user of the cluster.

In a typical parallel application cluster (Figure 4), clients communicate with clustered servers through a local area network (LAN). LANs are complex networks that allow several different types of devices to communicate with each other. Clients and servers are physically connected to the LAN cable by network interface controllers (NICs), which are placed inside the devices. Network applications manage the communications between the various devices on the LAN. These applications use the LAN NICs and multipoint protocols like *TCP/IP* to transmit messages between clients and servers.

TCP/IP - Transmission Control Protocol/Internet Protocol.

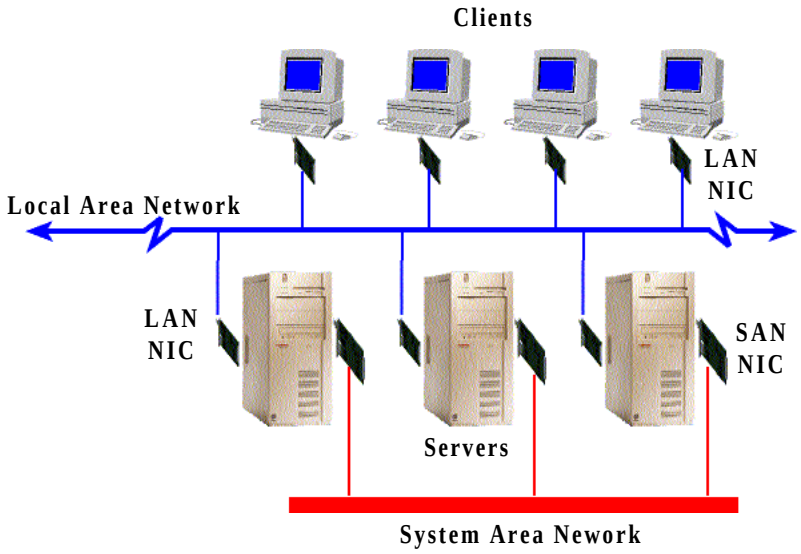

Figure 4. Representation of a LAN and a SAN in a parallel application cluster.

Clustered servers communicate with each other through a *system area network*, or SAN. Traditional network applications use LAN protocols like TCP/IP for server-to-server communication on the SAN. When an application sends messages through a SAN NIC, it uses physical memory. The application's request to use memory goes through the operating system (OS) kernel and is processed by the server's processor. This software overhead consumes processing resources, resulting in decreased processor utilization for productive tasks.

System area network – a high-performance network that interconnects servers and storage devices within a cluster.

A parallel application server cluster demands a SAN that can deliver high bandwidth, low latency, and vast expandability while consuming a negligible portion of the processing resources. Compaq, Intel, Microsoft, and over 100 industry contributors are meeting these demands with the Virtual Interface Architecture.

Virtual Interface Architecture for System Area Networks

Unlike a LAN, a SAN is a dedicated, highly reliable, and relatively secure computing environment. Therefore, high overhead protocols like TCP/IP are not required for clustered servers to communicate with each other through a SAN. Instead, applications on clustered servers can communicate using a distributed messaging interface called the Virtual Interface (VI) Architecture.

The VI Architecture is software that allows applications to send data directly to networking hardware, such as the SAN NIC, without going through the OS kernel (Figure 5).[2] The results are higher CPU utilization for productive tasks and higher overall performance within the cluster. Connection maintenance (control) functions still go through the kernel to access the NIC, but they do not occur during messaging, and therefore, do not affect latency.

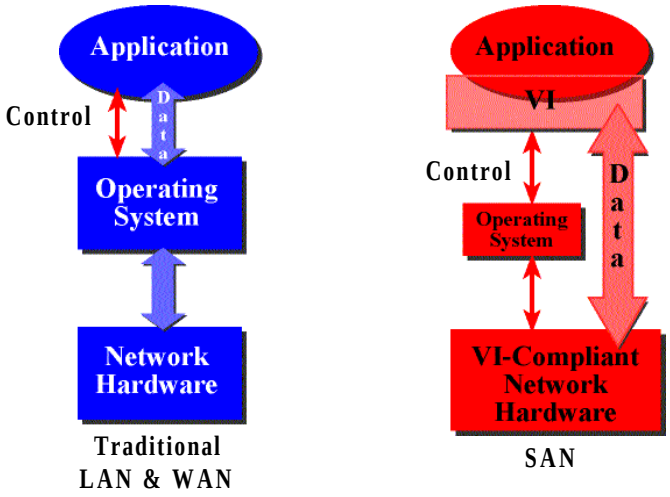

Figure 5. Compaq's VI Architecture implementation.

VI Architecture offers the following benefits:

- **Lower cost per transaction.** With VI Architecture as a broad industry standard, users will benefit from the lower cost of transactions and application systems. With VI Architecture on existing servers and storage devices, customers will not need to increase server capacity as they have in the past, thus reducing cost. Also, customers using traditional network interfaces for application corrections, will operate more efficiently in their use of buffer pools and networks, so overall communication costs will be reduced.
- **Interoperability.** Software vendors will be able to offer existing applications that take advantage of VI Architecture as Microsoft integrates the architecture in their Winsock 2 interface. This integration results in immediate interoperability for Winsock 2-enabled applications when VI Architecture suitable hardware is used in the systems interconnect. Similarly, other OS vendors can support VI Architecture under their communication mechanisms.
- **Scalability.** Reliable, high-bandwidth SANs allow the work traditionally done by protocol stacks to be offloaded to the VI Architecture-compliant NIC. This offloading allows the processor to spend more time doing useful work. Also, the low latency of the VI Architecture ensures that the processor will not be gated (throttled) by waiting on messaging. These features of VI Architecture greatly increase the scalability of parallel applications over traditional communications mechanisms.

Compaq, Intel, and Microsoft codeveloped VI Architecture as an industry-standard protocol and published it as a formal specification in December 1997. Over 100 companies have also contributed to the specification.[3] Compaq, Tandem, IBM, Oracle, and Giganet are producing hardware and software products that take advantage of the VI Architecture.

[2] See technology brief *Virtual Interface Architecture for SANs*, document number ECG058/0698, and white paper *Virtual Interface (VI) Architecture – The New Open Standard for Distributed Messaging Within a Cluster*, document number ECG098/0998, for more information.
[3] A complete list of contributors can be found on the VI Architecture web site at http://www.viarch.org.

ServerNet Technology

The VI Architecture-enabled SAN described above permits network interconnects to significantly reduce the overhead and processing requirements for communications between clustered servers. The Compaq E2000 Platform Architecture uses an industry-standard, highly reliable, high-speed, low-latency hardware interconnect called ServerNet to implement the VI Architecture. ServerNet was originally developed by Tandem in the early 1990's as a proprietary, high-speed interprocessor link for use with SANs. Due to ServerNet's success and Compaq's acquisition of Tandem, ServerNet is now being externalized to provide a universal, low-latency interconnect for processing, storage, and communications using a standards-based PCI adapter.

The ServerNet PCI adapter is designed with reliability and fault tolerance in mind. As shown in Figure 6, the adapter has two ports (X and Y). Each port can be connected to a separate ServerNet switch box for a redundant connection. If any single component on either port connection should fail, data will be transferred by the second connection, thus improving reliability. ServerNet also provides high reliability through self-checking built into the application-specific integrated circuits (ASICs) and hardware routing protocol. Hardware routing protocol ensures that end-to-end reliable data transfer occurs. [4]

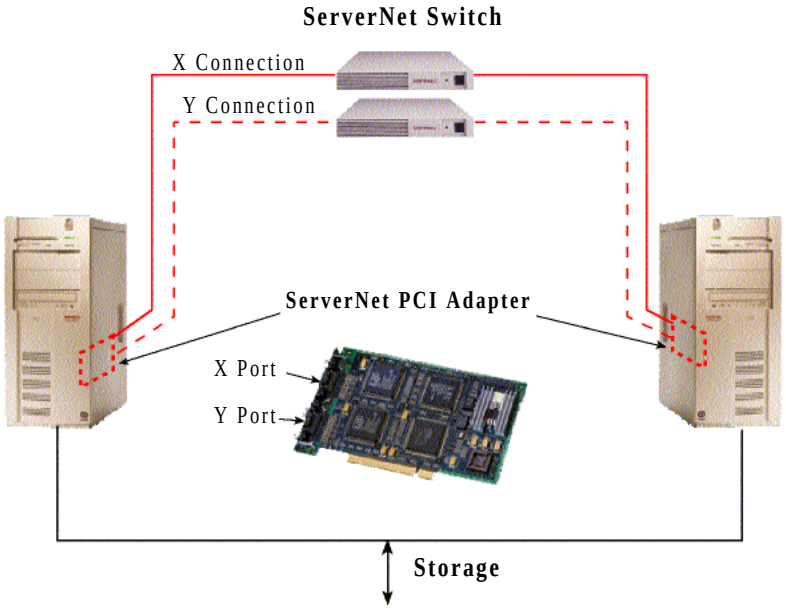

Figure 6: ServerNet PCI Adapter with dual ports.

Fabric – An active, intelligent, switching mechanism used to connect devices.

Instead of focusing on data processing as other communication interfaces do, ServerNet focuses on moving data. ServerNet considers all the major components of a SAN as independent elements so that any component—processor, storage, or I/O device—can directly communicate with any other component without processor intervention. ServerNet's any-to-any switching *fabric* enables virtually endless cluster growth. The fabric consists of many high-speed, point-to-point data paths. Each data path uses a hardware protocol to ensure that the data is delivered accurately. The data paths allow system elements (processors, storage, or I/O) to be joined into a cluster of servers. The data paths from the system elements are connected to the cluster by means of six-port ServerNet switches, which contain very large scale integration (VLSI) integrated circuit devices that route data to the correct paths.

[4] See technology brief *ServerNet Technology*, document number ECG043/0298, for more information.

A high-performance, scalable cluster requires two communication ingredients: high bandwidth and low latency. ServerNet scales I/O bandwidth by allowing multiple data paths to be added and cascaded to support as much bandwidth as needed. As more nodes are added to the cluster, more data paths are added, and the aggregate bandwidth of the ServerNet interconnect increases. When a data request is made in a ServerNet-based SAN, ServerNet uses the following techniques to ensure low latency:

- **Packet switching and wormhole routing.** ServerNet retrieves the data and divides it into 64-byte packets (Figure 7). The destination address is added to the front end of each packet, allowing the switch to route the packet to its final destination. If necessary, packets from the same data file can travel different paths to reach the destination. Additionally, as the first bytes of a packet reach the router, the router decodes the packet address and routes the head of the packet to its destination before the entire packet is received. This technique is called wormhole routing.
- **Push/pull approach.** When writing (pushing) data, at the beginning of the transaction, the source sends a message to the destination to allocate enough buffers to receive a larger message. The source waits for acknowledgment from the destination that the buffers are available before it sends all the data. When pulling (reading) data, the destination allocates buffers before it requests data. Then it transfers the data through the ServerNet adapters without interrupting the OS or application.

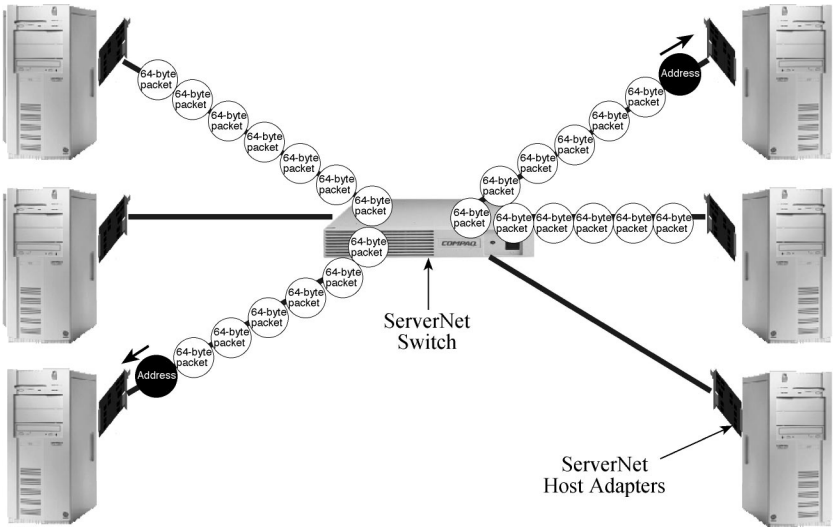

Figure 7. Wormhole routing of ServerNet data packets.

Gigabit Ethernet Technology

Bandwidth management – the best practices used by network designers and managers to balance traffic across networks.

The any-to-any network traffic enabled by technologies like ServerNet challenges traditional concepts of balancing networks using *bandwidth management*.[5] The 80-20 axiom, which held that 80 percent of network traffic remained in the LAN where it originated, has been invalidated by distributed computing and the explosive growth of Internet/intranet traffic. As a result, network congestion and reduced network performance are still issues for network managers.

The effects of network congestion stimulated the development of Ethernet (10 Mb/s), Fast Ethernet (100 Mb/s), and, subsequently, Gigabit Ethernet (1,000 Gb/s). Gigabit Ethernet employs the same *CSMA/CD* protocol, frame format, and frame size as its predecessors, plus interoperability and backward compatibility with installed Ethernet. Gigabit Ethernet is now used as a high-speed backbone interconnect between 10/100Base-T switches and as a connection to high-performance servers[6] (see Figure 8).

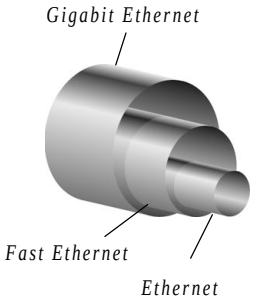

CSMA/CD (Carrier Sense Multiple Access with Collision Detection) - the combination of carrier sensing to prevent transmissions that otherwise would result in collisions and collision detection to terminate incomplete transmissions that have collided.

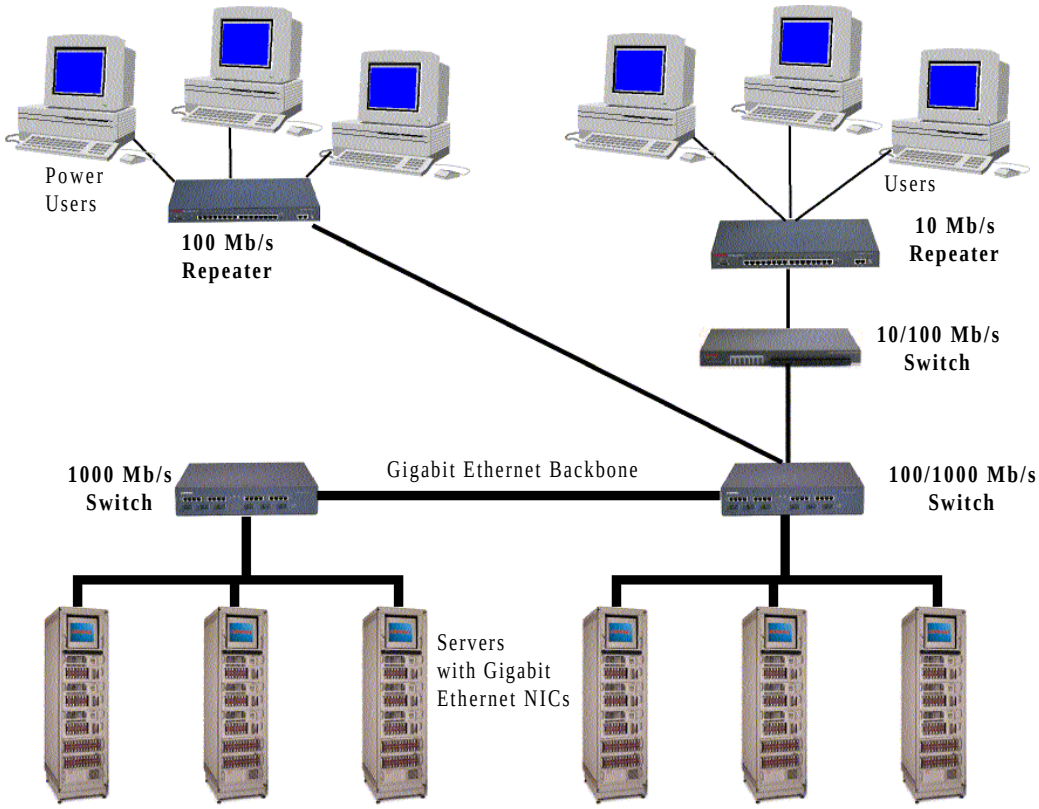

Figure 8. Gigabit Ethernet backbone in the E2000 Platform Architecture.

Gigabit Ethernet offers:

- Compatibility with existing applications, network OSs, and network management technologies.
- Support for full duplex operation for *point-to-point* connections and half-duplex operation for shared connections.

Point-to-point connection – a communications network that provides a path from one location to another.

[5] See white paper *Bandwidth Management and Gigabit Ethernet*, document number NAC125A/0498, for more information.
[6] See technology brief *Gigabit Ethernet Technology*, document number130A/1196, for more information.

Compaq is one of the founding members of the Gigabit Ethernet Alliance. It is an open forum whose purpose is to promote industry cooperation and to propose interoperability issues to the IEEE technical committee for the standardization of Gigabit Ethernet.

Fibre Channel Technology

Channel – a direct or switched point-to-point connection that transfers data at the highest possible speed.

Fibre Channel is the general name of an integrated set of standards being developed by committees accredited by the American National Standards Institute (ANSI). Compaq is a member of the Fibre Channel ANSI committees and is actively working to enhance the standards. Fibre Channel is an industry-standard interconnect. Fibre Channel is a high-performance serial I/O interface that is media independent and supports simultaneous transfer of many different protocols. It uses fiber to transfer data at the fastest speed possible and contains the features to provide extensive connectivity, distance, and protocol multiplexing.

Each Fibre Channel device (node) has one or more ports for external communication. Each port uses two fibers, one to send data and one to receive data. The pair of fibers is called a link. All the components that connect ports comprise an interconnect topology.

Various topologies are used to connect Fibre Channel ports. The two basic Fibre Channel topologies used today are the fabric switch, which has the highest performance, and the arbitrated loop (see Figure 9). Every port has a unique address—analogous to a telephone number—that allows it to communicate directly with another port.

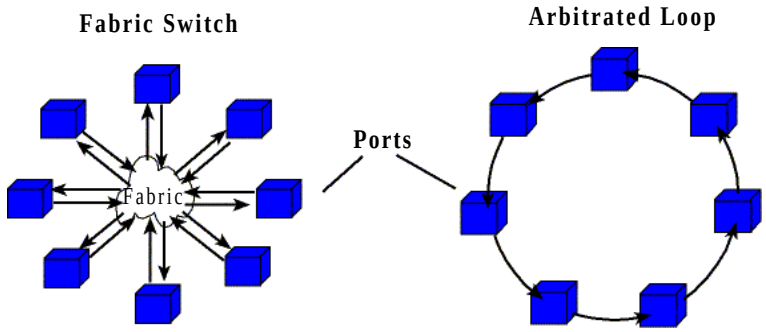

Figure 9. Fabric switch and Fibre Channel Arbitrated Loop topologies.

The fabric switch, which is analogous to a telephone switching system, uses the unique node addresses to simultaneously connect multiple pairs of nodes. Each node (device) port is connected to a port on the fabric switch just as a phone, fax, or modem is plugged into a phone jack. The internal structure of the fabric is transparent to the node ports, so the fabric appears as a single interconnecting entity. The throughput capacity of the switched topology can increase incrementally as more nodes are added. The size of the fabric can grow to be very large, constrained by the addressing limit of 2^{24} (more than 16 million addresses) for a single fabric.

The Fibre Channel Arbitrated Loop (FC-AL) is a serial interface that creates logical point-to-point connections between ports with the minimum number of transceivers; therefore, it provides a lower cost solution. The bandwidth of an FC-AL is shared by all ports on the loop. Only a single pair of ports on the loop can communicate at one time, while the other ports on the loop act as repeaters.

Fibre Channel Hubs are used to facilitate the wiring of an FC-AL. A hub contains several ports that are internally connected in a loop (Figure 10). Each port is fitted with a port bypass switch to maintain the continuity of the loop should a controller or device attached to the port be powered off or malfunction. Hubs can be cascaded to provide additional ports for more connectivity. More than 100 nodes can be attached on a single FC-AL.

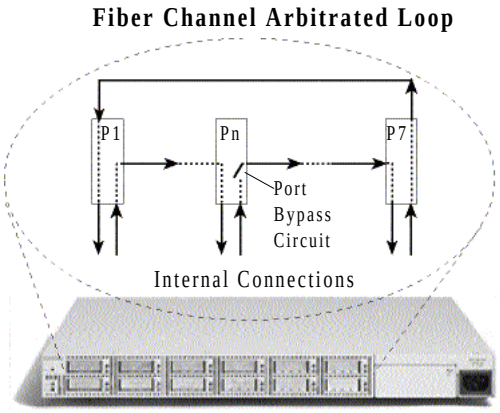

Figure 10. The internal connections of a Compaq Storage Hub 7 create a Fibre Channel Arbitrated Loop.

Fibre Channel provides scalability in bandwidth, distance, and connectivity.

- *Bandwidth* - Fibre Channel bandwidth, currently at 100 MB/s per link in both directions, is targeted to grow to 200 MB/s and 400 MB/s per link in the future.
- *Distance* - Fibre Channel can transfer data for distances up to 500 meters using 50-micron multimode fibre and up to 10 kilometers using 9-micron single mode fibre. This capability allows remote mirroring of data at very high data rates and allows multiple systems to share the same storage. Work is in progress to extend the distance even further.
- *Connectivity* - Fibre Channel provides greater connectivity per PCI slot, and thus greater capacity. Large subsystems of devices can be attached to ports on hubs, and expansion modules can be cascaded to connect far more subsystems. The storage capacity of a single I/O slot can be scaled to more than 1.6 TB (more than 9.6 TB for a six-PCI slot server) with the current Compaq products.

Compaq's strategic direction for high-performance and high-capacity external storage is based on Fibre Channel technology (see Figure 11). Fibre Channel is a key technology for the high-speed storage interconnect (that is, processor-to-storage and storage-to-storage communications) and for the serial drive interface (high-performance disk systems). It allows the integration of primary and secondary storage, as well as allowing shared storage among multiple servers.[7]

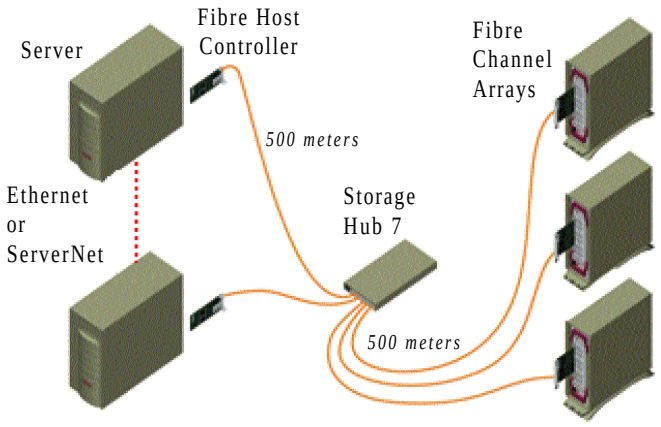

Figure 11. Fibre Channel Storage System Cluster Configuration.

[7] For more information, see *Strategic Direction for Compaq Fibre Channel-Attached Storage*, document number ECG009/1097.

Fibre Channel, Gigabit Ethernet, and ServerNet have a common architecture at the physical level, which promotes the use of common parts and allows the use of the same infrastructure. Gigabit Ethernet leverages the physical level and the encoding used in Fibre Channel. Gigabit Ethernet provides a high-speed local area network, while Fibre Channel provides a high-speed storage area network. A Fibre Channel storage area network allows a client attached to a specific processor to access data in any storage device within the storage area network because all storage devices are accessible to all processors.

In addition, Compaq is developing ServerNet as the node-to-node interconnect within Compaq clusters because ServerNet features very low latency. Future implementations of ServerNet will use fiber and leverage the physical level and encoding of Fibre Channel.

HIGH-AVAILABILITY TECHNOLOGIES

The E2000 Platform Architecture provides integrated solutions from Compaq and its partners to meet all levels of availability from high-availability, single-server solutions to NonStop® continuous-availability cluster solutions. To increase availability at the server level, Compaq has implemented several technologies, including:

- Redundant network interface controller technology
- Redundant, hot-plug fans
- Redundant processor power modules
- Redundant, hot-plug power supplies
- Hot-plug drives
- Compaq Smart 3100ES Array Controller
- PCI Hot Plug technology for PCI adapters

At the cluster level, the E2000 Platform Architecture includes solutions ranging from high-availability 2-node failover clusters up to continuous availability n-node parallel application clusters and SANs. Microsoft Cluster Server, Oracle Parallel Server, and the integration of NonStop technologies from Tandem allow both scalability and availability.

These server level and cluster level high-availability technologies are described in the following sections.

High-Availability Server Technologies

Redundant Network Interface Controller Technology

Compaq's Redundant Netelligent Network Interface Controller (NIC) technology allows two similar NICs to share a single instance of device driver code. One NIC becomes the active network controller and the other NIC acts as a standby controller. If the active NIC fails, the network traffic can be switched automatically to the standby NIC. This redundancy eliminates the NIC or cable as a single point of failure. With PCI Hot Plug technology, the failed NIC can be replaced without shutting down the system. Thus, the end-user can have continuous service and the administrator can greatly reduce both planned and unplanned downtime.[8]

[8] See white paper *High Availability Options Supported by Compaq Network Interface Controllers*, document number NSC 209A/0698, for more information.

PCI Hot Plug Technology

The Compaq ProLiant 7000 and ProLiant 6500 servers are the first standards-based servers to incorporate industry-standard PCI Hot Plug technology. This technology, pioneered by Compaq, adds hot-plug capabilities to existing PCI adapters for increased system availability and serviceability. Compaq led the industry by licensing its PCI Hot Plug implementation to Intel, thus broadening customer access to this standard technology.

Compaq's implementation of PCI Hot Plug differs from implementations by other vendors whereby several slots are powered down at once or an entire PCI bus is powered down. Compaq incorporated electronics into the server to control each PCI Hot Plug slot individually, so that a single slot can be powered down without affecting the operation of the other slots.

Compaq's implementation of the PCI Hot Plug standard allows the following capabilities while the system is running:

- *Hot replacement*—replacing a single PCI adapter with a similar adapter
- *Hot upgrade*—replacing a PCI adapter with an upgraded adapter
- *Hot expansion*—adding a PCI adapter to an empty slot

The ProLiant 7000 is designed to support all three hot-plug capabilities for all 64-bit and 32-bit PCI Hot Plug slots. These capabilities can be added incrementally, depending on the OS implementation.

For more information about major OS vendors that support PCI Hot Plug, refer to the following technology briefs:

- *PCI Hot Plug Technology*, document number ECG080/0698
- *PCI Hot Plug Technology with Novell Architecture*, document number ECG081/0698
- *PCI Hot Plug Technology with SCO Software Architecture*, document number ECG082/0698
- *Deploying PCI Hot Plug on Compaq Servers in a Microsoft Windows NT Environment*, document number 064A/0797

PCI Hot Plug technology is backward compatible with existing industry-standard adapters. While new drivers are required to implement hot-plug capabilities, no changes are required to the vast majority of existing adapters. These leading independent hardware suppliers are committed to modifying their device drivers to be hot-plug aware: Adaptec, AMI, Dialogic, Digi International, Madge, Mylex, QLogic, Standard Microsystems Corporation, SysKonnect, and 3Com. In addition, Compaq is delivering hot-plug capable drivers for its own leading PCI adapters.

PCI Hot Plug technology adds to other high-availability solutions from Compaq such as Recovery Server Options, hot-pluggable SCSI drives, and redundant hot-pluggable power supplies, fans, and processor power modules.[9] Redundant components ensure that the server is highly fault tolerant. In many cases, if a power supply, fan, processor power module, or I/O board fails, the redundant standby component can take over operation with no downtime to the server. With the addition of hot-plug capabilities, server downtime can be eliminated while components are being replaced, added, or upgraded.

[9] See technology brief *Compaq ProLiant 7000 Server Technology*, document number ECG078/0798, for more information.

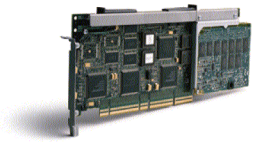

Smart Array 3100ES Controller

Compaq Smart Array 3100ES Controller

The Compaq Smart Array 3100ES Controller[10] is a 3-channel, Wide-Ultra SCSI-3 controller that ships standard in the ProLiant 7000. The architecture of the Smart Array 3100ES improves on the Compaq SMART-2 Array Controller by:

- Simplifying SCSI cabling configuration. The extended SCSI connector routes three Wide-Ultra SCSI-3 buses to the server's three internal drive cages.
- Improving serviceability through PCI Hot Plug capabilities.
- Increasing performance with 64 MB of removable, battery-backed cache (56 MB for read/write cache).
- Allowing all three drive cages to be configured as one contiguous redundant array of independent disks (RAID).
- Providing faster memory.
- Providing future support for controller redundancy.

High-Availability and Continuous Availability Cluster Technologies

Compaq cluster solutions range from high-availability 2-node failover clusters up to continuous availability n-node parallel application clusters. The Compaq ProLiant Cluster Series S is a 2-node high-availability solution for customers who wish to continue using Compaq's SCSI-based, industry-standard servers and ProLiant Storage Systems. The Compaq ProLiant Cluster Series S Model 100 (Figure 12) also uses SMART-2 Array Controllers, Recovery Server Option (RSO), and a ProLiant Cluster/S 100 Kit. The kit contains the integrated software and documentation needed to provide a high-availability cluster for applications and data in a business-critical environment.[11]

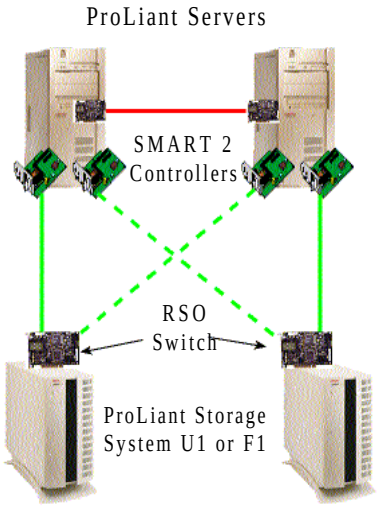

Figure 12. ProLiant Series Cluster Series S

[10] See technology brief *Compaq Smart Array Controller Technology*, document number ECG079/0698, for more information.
[11] See white paper *Order and Configuration Guide for Compaq ProLiant Cluster Series S Model 100*, document number ECG029/0498, for more information.

The Compaq ProLiant Cluster Series F has even higher levels of performance and availability (Figure 13 Left). It uses Compaq's industry-standard servers, Ethernet or Tandem ServerNet interconnect, Compaq Fibre Channel Storage Systems, and Compaq installation and management utilities.[12] The ProLiant Cluster/F 100 kit gives customers tremendous flexibility in how they build their cluster. For example, if a customer already has a single ProLiant 2500 server and needs a higher level of availability, he can purchase the cluster kit and a ProLiant 6500 as a second node to create a high-availability cluster.

The Compaq Parallel Database Cluster Model PDC/O1000 (Figure 13 right) provides high availability and high scalability of Oracle database environments using industry-standard components and software. The Compaq PDC/O1000 cluster is certified for Oracle 8 Server Release 8.0.5 and Oracle Parallel Server. It combines fault-resilient, rack-mounted Compaq servers, 100Base-T Ethernet IPC, and Compaq Fibre Channel Storage with Oracle Parallel Server and Oracle 8 Release 8.0.5 relational database management software. The cluster acts as a single logical server for a single Oracle database. Users can concurrently access the clustered database on read and write operations. Each node of the cluster can work independently of the others, allowing related client/server applications to access the same set of data. If any server in the cluster fails, the remaining nodes can continue to access the database, ensuring high availability for business-critical applications.

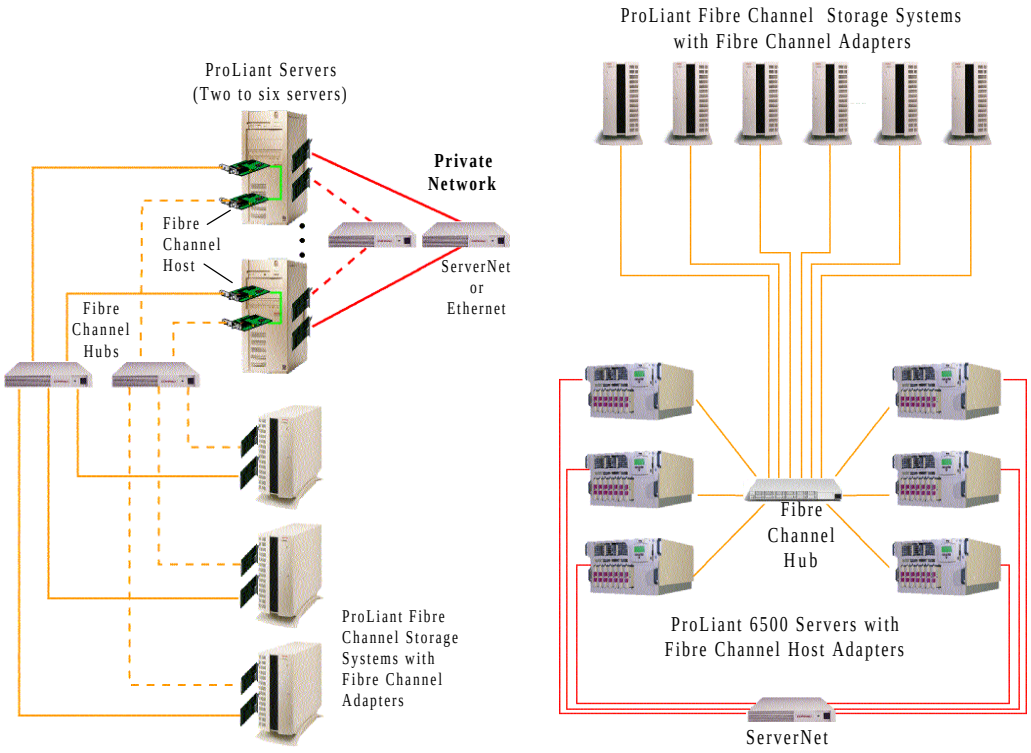

Compaq ProLiant Cluster Series F **Oracle Parallel Database Cluster**

Figure 13. Compaq ProLiant Cluster solutions provide high availability in business-critical environments.

[12] See white paper *Order and Configuration Guide for Compaq ProLiant Cluster Series F Model 100*, document number ECG030/0998, for more information.

For some users, the need for uninterrupted service has grown beyond high availability to "continuous availability." Continuous availability requires the elimination of both planned and unplanned downtime. Parallel application clusters offer continuous availability by allowing other servers to pick up the workload of a failed server. To users, the applications and data running on the failed server remain available.

The Tandem NonStop® Clusters for SCO UnixWare® (NCU) platform is designed from the ground up with continuous availability in mind.[13] It offers users the ability to run all existing applications without modification and provides a fault-tolerant environment to perform critical business. It also provides an easy migration path for scaling to more powerful hardware and software without disruption of service to users. Figure 14 shows a representation of the Tandem NCU, which consists of:

Middleware – software that sits between an application and a control program (operating system, network control program, database management system, etc.) that provides a single programming interface for application development.

- SCO UnixWare 2.1.2
- Tandem NCU Clustering *Middleware*
- ServerNet PCI Cards and cabling
- ServerNet 6-node Switch (Hub)
- Two to six Compaq ProLiant Servers: ProLiant 850R, ProLiant 2500, or ProLiant 6500
- Integrity XC clusters turnkey packaging
- Compaq ProLiant Storage System U
- Recovery Server Option (RSO)

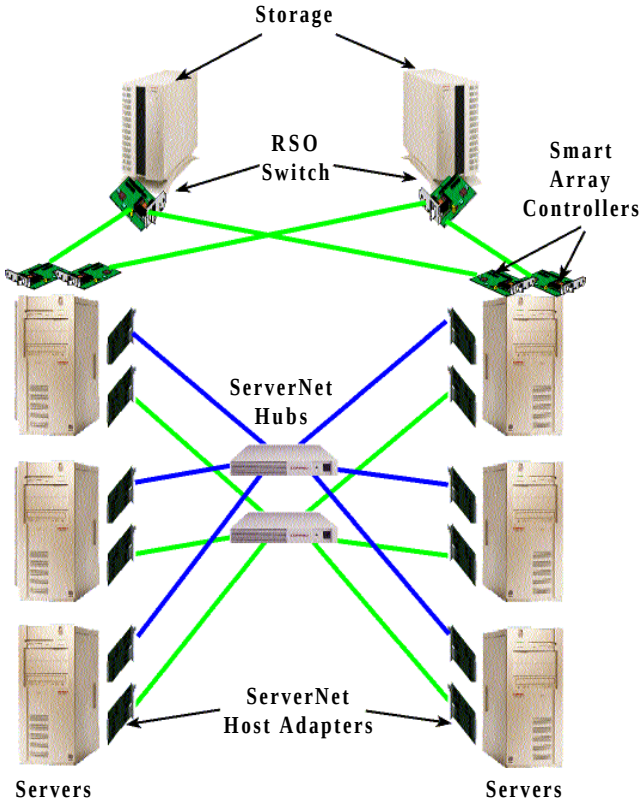

Figure 14. NonStop® Clusters for SCO UnixWawre®.

[13] See technology brief *Tandem's NonStop Clusters for SCO UnixWare Provides a Fault-Tolerant Operating Environment for Clustered Computers* at www.tandem.com.

MANAGEABILITY

In addition to the high-availability and high-scalability technologies mentioned previously, the Compaq E2000 Platform Architecture delivers the right combination of management and integration tools required to optimize and maintain distributed computing environments. Manageability measures the ease with which IT staff can administer and troubleshoot all of the components of the IT structure. The E2000 Platform Architecture fully embraces standards-based management tools, including:

- Compaq Insight Manager
- Integrated Management Display
- Integrated Remote Console
- Web-Based Enterprise Management (WBEM)

Compaq Insight Manager

Compaq Insight Manager is an intelligent management tool that constantly monitors and analyzes critical information about every Compaq server on the network. Compaq Insight Manager brings this information to the customer and gives them the tools to respond. Customers can quickly and easily address problems before they become disasters, thus preventing downtime and costly data loss. Compaq Insight Manager is based on the 32-bit architecture, and it is powerful enough to manage thousands of systems from a single console.

Compaq Insight Manager supports the integration management process through the Integration Server Maintenance feature, which lets customers compare their current software versions with the latest software available from Compaq, and then recommends needed updates. Customers can download the updates via the Internet and manage upgrades to meet their needs.[14]

Integrated Management Display

The Integrated Management Display (IMD), shown in Figure 15, is a backlit liquid crystal display (LCD) that is menu-driven and allows administrators to enter server and contact information. Just as importantly, the IMD displays critical information, warning messages, and error messages in an easy-to-understand format.

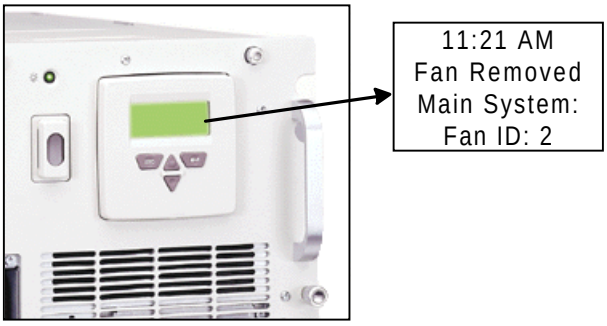

Figure 15: Integrated Management Display with sample alert message.

A 5V auxiliary power line supplies power to the IMD whenever the server power line is plugged in. Even if the server power goes down, the auxiliary line continues to power the IMD so that critical information, such as the last error message and the administrator's name, can be obtained.

[14] See white paper *Compaq Insight Manager 4.2*, document number ECG044/0998 for more information.

The IMD unit contains its own static memory (64 K x 8 SRAM). This memory contains both the software code to control the IMD and the text messages shown on the display. When the server initiates power-on self-test (POST), the code downloads from the system ROM. Therefore, any software enhancements added in the future will be available through a flashable ROMpaq. Also during POST, all event logs, system information, and administrator information upload from the system non-volatile RAM. During runtime, new events are stored in both the non-volatile RAM and the IMD SRAM. If the IMD unit is removed and replaced, all key event and system information is still available.

Integrated Remote Console

Integrated Remote Console (IRC) consists of a Compaq-designed application specific integrated circuit (ASIC) and associated firmware. IRC allows out-of-band, or asynchronous, management capabilities such as remote console and remote reset. These capabilities are independent of the state of the network OS. With IRC, an administrator has full text mode video and keyboard access to a server, even if the OS is down. The administrator has the ability to access the server, perform diagnostics, reset the system, watch the reset process remotely, and view reset sequences, regardless of whether the server OS is functional.

For more information about IRC, refer to the Compaq technology brief *Remote Server Management with Integrated Remote Console,* document number 582A/1096.

Web-Based Enterprise Management Technology

Web-based enterprise management (WBEM) is an ongoing industry initiative to integrate current enterprise management technology with the latest advances in web technology, providing organizations with an easy-to-use, cost-effective method for managing enterprise systems. WBEM shifts the focus from the infrastructure of network management to the content and processes. Administrators require the ability to track network problems from the user or application level through the system and various network layers to the remote service or servers. Based on industry standards, the new WBEM architecture combines existing and new management elements and presents them through web browsers.[15] Instead of a two-tier architecture (management console and managed device), WBEM uses a three-tier architecture that includes a management server, managed devices, and a web browser. Figure 16 illustrates the WBEM architecture.

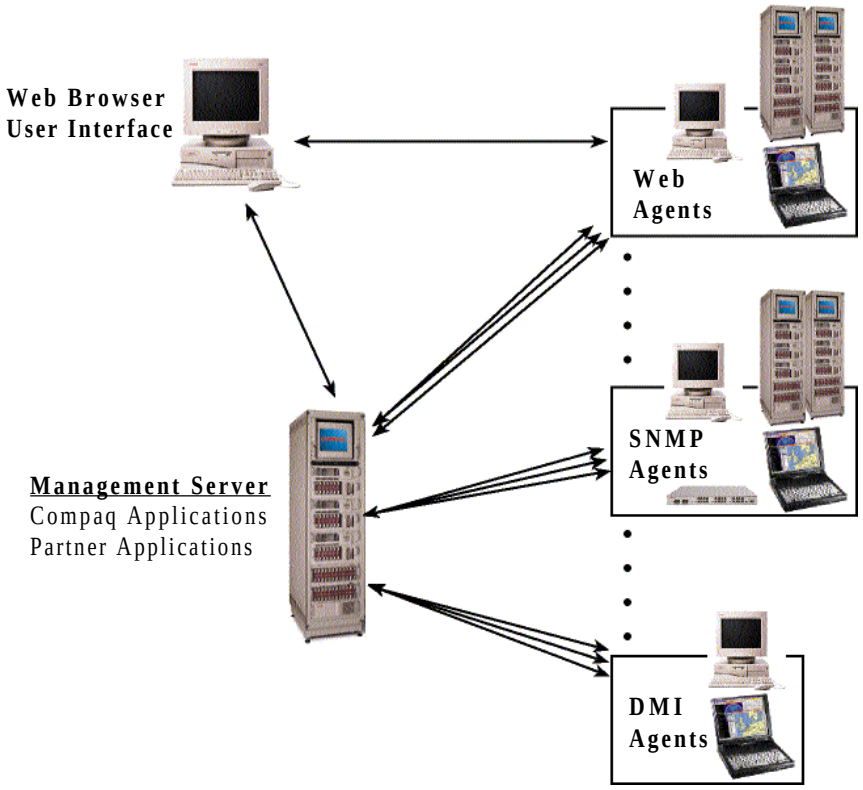

Figure 16: WBEM architecture model. A browser uses web-type connections to communicate with the server or managed elements over an intranet.

In this new architecture, the web browser replaces the standard management console, providing access to either the web agents or the management server. The management server facilitates the web access to web-enabled, simple network management protocol (SNMP), or desktop management interface (DMI) agents. The management server also performs the functions required to logically aggregate multiple managed objects sharing some common action, fault, or configuration of interest. Self-describing web agents use common web-enabling components to provide registration, discovery, security, HTTP communications, and a home page. The agents can render information in HTML for viewing by a browser, or they can send information directly to the management server.

[15] See technology brief *Web-Based Enterprise Management—An Ongoing Initiative*, document number ECG143/0398, for more information.

CONCLUSION

The Compaq E2000 Platform Architecture provides enterprise customers with powerful, reliable, and flexible distributed computing systems to run mission-critical applications. The architecture is based on standards-based components, including SMP servers, high-performance networking technologies, and high-capacity storage systems, connected by scalable system area network interconnects. Using these industry-standard building blocks, customers can create highly scalable clustered server systems that offer high availability to continuous availability.

High Scalability

The E2000 Platform Architecture allows scalability in all aspects of enterprise computing systems. VI Architecture-enabled SANs improve the performance of clustered SMP servers, which results in a lower cost per transaction and higher scalability for parallel application clusters. ServerNet, a reliable, high-speed, low-latency hardware interconnect, allows these powerful clusters to be connected with robust Fibre Channel storage, which can be scaled to more than 200 TB. These systems will be tightly integrated with industry-standard operating systems like Microsoft Windows NT and SCO UnixWare and with enterprise-caliber applications from Microsoft, Oracle, SAP, and other independent software vendors.

High Availability

Compaq has formed strategic alliances with industry partners to meet customer demands for highly available, industry-standard, enterprise computing systems. At the server level, Compaq provides redundant controllers, fans, processor power modules, and power supplies as well as PCI Hot Plug drives to greatly reduce planned and unplanned downtime. At the cluster level, solutions range from highly-available 2-node failover clusters to continuously available n-node parallel application clusters and SANs.

Manageability

The E2000 Platform Architecture delivers standards-based management and integration tools to optimize and maintain distributed enterprise computing environments. From intelligent management tools that constantly monitor and analyze the status of every server, to web-based technology that provides automated management of networked enterprise systems, the E2000 Platform Architecture makes it easier to administer and troubleshoot distributed resources.

The Bottom Line

The bottom line is that the standards-based solutions in the Compaq E2000 Platform Architecture will allow enterprise customers to scale processing, storage, and communications independently and provide them with a wider choice of systems, applications, and services at a lower cost.